Cisco Secure Firewall Management Center Device Configuration Guide, 7.7

First Published: 2025-03-05

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2025 Cisco Systems, Inc. All rights reserved.

CONTENTS

PART I **Getting Started with Device Configuration** 83

CHAPTER 1 **Device Management** 1

About Device Management 1

About the Firewall Management Center and Device Management 1

What Can Be Managed by a Secure Firewall Management Center? 2

About the Management Connection 3

Beyond Policies and Events 3

About Device Management Interfaces 4

Management and Event Interfaces on the Firewall Threat Defense 4

Using the Firewall Threat Defense Data Interface for Management 4

Management Interface Support Per Device Model 5

Network Routes on Device Management Interfaces 6

NAT Environments 7

Management and Event Traffic Channel Examples 9

Requirements and Prerequisites for Device Management 10

Log Into the Command-Line Interface on the Device 11

Complete the Firewall Threat Defense Initial Configuration for Manual Registration 13

Complete the Firewall Threat Defense Initial Configuration Using the Firewall Device Manager 13

Complete the Firewall Threat Defense Initial Configuration Using the CLI 19

Configure an Event Interface 26

Manage Devices 28

About the Device Management Page 28

Add a Device Group 30

Register With the Management Center 31

Registration Key Method 31

Serial Number Method (Zero-Touch Provisioning) 41
Add a Chassis 53
Register With a New Management Center 56
Resolve Serial Number (Zero-Touch Provisioning) Registration Issues 57
Unregister a Device from the Firewall Management Center 60
Shut Down or Restart the Device 62
Download the Managed Device List 63
Migrate Firewall Threat Defense Devices 63
Supported Devices for Migration 63
License for Migration 65
Prerequisites for Migration 65
What Configurations Does the Wizard Migrate? 66
Guidelines and Limitations for Migration 68
Migrate a Secure Firewall Threat Defense 69
Best Practices for Threat Defense Device Migration 71
Switch Managers 72
Switch from the Firewall Device Manager to the Firewall Management Center 72
Switch from Firewall Management Center to Firewall Device Manager 77
Hot Swap an SSD on the Secure Firewall 3100/4200 78
Disable the USB Port 81
Disable the USB Port on a Device 81
Disable the USB Port in Multi-Instance Mode 82
History for Device Management 84

CHAPTER 2 **Device Management Using Device Templates** 89
About Device Management using Device Templates 89
Methods to Register Devices using Templates 90
Variables and Network Object Overrides 90
Model Mapping 90
Templates and High Availablity 90
Templates and Domains 91
Validation of Template Configuration Before and After Application of Template on Device 92
Requirements and Prerequisites for Device Management using Device Templates 93
Licenses for Device Management using Device Templates 94

Guidelines and Limitations for Device Management using Device Templates 94
Workflow for Device Management using Device Templates 98
Configure Device Templates 98
Add a Device Template 98
Create a New Device Template 98
Generate a New Device Template from an Existing Device 99
Import a Device Template 99
Configure Device Settings in the Template 100
Add a Physical Interface 100
Add a Logical Interface 101
Edit an Interface 101
Configure Other Device Settings 102
Configure Template Settings 102
Edit General Settings 103
Edit Licenses 103
Edit Applied Policies 104
Edit Advanced Settings 104
Edit Deployment Settings 105
Configure Template Parameters 106
Add Model Mapping 109
Configure Site-to-Site VPN Connections in a Device Template 110
Device Template with Site-to-Site VPN Connections Workflow 110
Configure an SD-WAN VPN Connection 110
Configure a Route-Based Site-to-Site VPN Connection 111
Configure a Policy-Based Site-to-Site VPN Connection 113
Register a Device Using Device Template and Add it to a Route-Based VPN Topology 114
Add a Device to an SD-WAN Topology in a Dual ISP Deployment 115
Configure a Template for Firewall Threat Defense Devices Managed Using the Data Interface 117
Use Templates with Devices 119
Add a Device to the Management Center Using a Device Template 119
Apply Templates to Existing Devices 119
Apply a Template 119
Reapply a Template 120
Monitoring Device Templates 120

View Associated Devices 121
Generate a Template Apply Report 121
Audit Logs 122
Troubleshooting Device Templates 122
History for Device Management using Device Templates 125

CHAPTER 3 **Device Settings** 127
Edit General Settings 127
Generate Troubleshooting Files 129
View CLI Output 132
Copy a Configuration to Another Device 134
Export and Import the Device Configuration 136
Edit License Settings 141
View System Information 141
View Device Inventory 142
View the Inspection Engine 142
Edit Health Settings 143
Out-of-Band Configuration Detection 143
Guidelines for Out-of-Band Configuration 144
Access Recovery-Config Mode in the Diagnostic CLI 146
Acknowledge the Out-of-Band Configuration 149
Edit Management Settings 153
Configure a Redundant Manager Access Data Interface 153
Change Manager Access Interface Settings 158
Change the Device IP Address 158
Change the Firewall Management Center IP Address 172
Change Both Firewall Management Center and Threat Defense IP Addresses 173
Change the Manager Access Interface 177
Troubleshooting the Management Connection 189
Manually Roll Back the Configuration if the Firewall Management Center Loses Connectivity 189
Troubleshoot Management Connectivity on a Data Interface 190
View Inventory Details 195
Edit Applied Policies 196
Edit Advanced Settings 197

Configure Automatic Application Bypass 198
Configure Object Group Search 199
Configure Interface Object Optimization 200
Edit Deployment Settings 201
Edit Cluster Health Monitor Settings 204
History for Device Settings 209

CHAPTER 4 Users 213
About Users 213
Internal and External Users 213
CLI Access 213
CLI User Roles 214
Requirements and Prerequisites for User Accounts for Devices 214
Guidelines and Limitations for User Accounts for Devices 215
Add an Internal User at the CLI 215
Configure External Authentication for the Firewall Threat Defense 217
About External Authentication for the Firewall Threat Defense 217
About LDAP 218
About RADIUS 218
Add an LDAP External Authentication Object for Firewall Threat Defense 219
Add a RADIUS External Authentication Object for Firewall Threat Defense 224
Enable External Authentication for Users on Firewall Threat Defense Devices 229
Troubleshooting LDAP Authentication Connections 229
History for Users 231

CHAPTER 5 Change Management 233
About Change Management 233
How to Configure Devices in the Change Management Workflow 234
Creating Separate Approver and Configuration Roles 234
Policies and Objects that Support Change Management 235
Requirements and Prerequisites for Change Management 237
Guidelines and Limitations for Change Management 237
Enabling or Disabling Change Management 238
Managing Tickets 239

Creating Change Management Tickets 240
Opening a Ticket for Configuration Changes 241
Previewing a Ticket 241
Submitting a Ticket 242
Discarding a Ticket 242
Approving or Rejecting a Ticket 243
Taking Over or Reassigning Tickets 243
History for Change Management 244

CHAPTER 6 **Configuration Deployment** 247
About Configuration Deployment 247
Configuration Changes that Require Deployment 247
Deployment Preview 248
Selective Policy Deployment 249
System Username 251
Auto-Enabling of Application Detectors 252
Asset Rediscovery with Network Discovery Policy Changes 252
Snort Restart Scenarios 252
Restart Warnings for Devices 252
Inspect Traffic During Policy Apply 253
Snort Restart Traffic Behavior 254
Configurations that Restart the Snort Process When Deployed or Activated 256
Changes that Immediately Restart the Snort Process 257
Requirements and Prerequisites for Policy Management 257
Best Practices for Deploying Configuration Changes 258
Deploy the Configuration 259
Deploy Configuration Changes 259
Redeploy Existing Configurations to a Device 265
Manage Deployments 266
View Deployment Status 266
View Deployment History 267
Set the Number of Configuration Versions 271
Roll Back a Deployment 272
Perform a Rollback 274

View the Deployment Rollback Transcript 276
Download Policy Changes Report for Multiple Devices 276
Compare Policies 277
Generate Current Policy Reports 278
History for Configuration Deployment 279

PART II **Device Operations** 283

CHAPTER 7 **Transparent or Routed Firewall Mode** 285
About the Firewall Mode 285
About Routed Firewall Mode 285
About Transparent Firewall Mode 286
Using the Transparent Firewall in Your Network 286
Passing Traffic For Routed-Mode Features 286
About Bridge Groups 287
Bridge Virtual Interface (BVI) 287
Bridge Groups in Transparent Firewall Mode 287
Bridge Groups in Routed Firewall Mode 288
Allowing Layer 3 Traffic 289
Allowed MAC Addresses 289
BPDU Handling 289
MAC Address vs. Route Lookups 290
Unsupported Features for Bridge Groups in Transparent Mode 291
Unsupported Features for Bridge Groups in Routed Mode 292
Default Settings 293
Guidelines for Firewall Mode 293
Set the Firewall Mode 294

CHAPTER 8 **Logical Devices on the Firepower 4100/9300** 297
About Interfaces 297
Chassis Management Interface 297
Interface Types 298
FXOS Interfaces vs. Application Interfaces 300
Shared Interface Scalability 302

Shared Interface Best Practices 302
Shared Interface Usage Examples 304
Viewing Shared Interface Resources 310
Inline Set Link State Propagation for the Firewall Threat Defense 311
About Logical Devices 311
Standalone and Clustered Logical Devices 312
Logical Device Application Instances: Container and Native 312
Container Instance Interfaces 312
How the Chassis Classifies Packets 313
Classification Examples 313
Cascading Container Instances 317
Typical Multi-Instance Deployment 318
Automatic MAC Addresses for Container Instance Interfaces 319
Container Instance Resource Management 320
Performance Scaling Factor for Multi-Instance Capability 320
Container Instances and High Availability 320
Container Instances and Clustering 320
Licenses for Container Instances 320
Requirements and Prerequisites for Logical Devices 321
Requirements and Prerequisites for Hardware and Software Combinations 321
Requirements and Prerequisites for Container Instances 323
Requirements and Prerequisites for High Availability 324
Requirements and Prerequisites for Clustering 325
Guidelines and Limitations for Logical Devices 328
Guidelines and Limitations for Interfaces 328
General Guidelines and Limitations 331
Configure Interfaces 332
Enable or Disable an Interface 332
Configure a Physical Interface 332
Add an EtherChannel (Port Channel) 333
Add a VLAN Subinterface for Container Instances 335
Configure Logical Devices 336
Add a Resource Profile for Container Instances 337
Add a Standalone Firewall Threat Defense 338

Add a High Availability Pair 343
Change an Interface on a Firewall Threat Defense Logical Device 344
Connect to the Console of the Application 346
History for Logical Devices 348

CHAPTER 9 **Multi-Instance Mode for the Secure Firewall 3100/4200** 355
About Multi-Instance Mode 355
Multi-Instance Mode vs. Appliance Mode 355
Chassis Management Interface 356
Instance Interfaces 356
Interface Types 357
Chassis Interfaces vs. Instance Interfaces 357
Shared Interface Scalability 359
Shared Interface Best Practices 359
How the Chassis Classifies Packets 361
Classification Examples 362
Cascading Instances 365
Typical Multi-Instance Deployment 366
Automatic MAC Addresses for Instance Interfaces 367
Performance Scaling Factor for Multi-Instance Mode 368
Instances and High Availability 368
Licenses for Instances 368
Requirements and Prerequisites for Instances 368
Guidelines and Limitations for Instances 370
Configure Instances 372
Convert a Device to Multi-Instance Mode 372
Configure Chassis Interfaces 376
Configure a Physical Interface 376
Configure an EtherChannel 379
Configure a Subinterface 383
Add an Instance 385
Customize the System Configuration 391
Configure SNMP 391
Import or Export the Chassis Configuration 392

Configure Chassis Platform Settings 395
Create a Chassis Platform Settings Policy 395
Configure DNS 396
Configure SSH and SSH Access List 397
Configure Syslog 402
Configure Time Synchronization 407
Configure Time Zones 409
Manage Multi-Instance Mode 410
Enable Multi-Instance Mode at the CLI 410
Change Interfaces Assigned to an Instance 413
Change Chassis Management Settings at the FXOS CLI 415
Monitoring Multi-Instance Mode 418
Monitoring Multi-Instance Setup 418
Monitoring Instance Interfaces 419
History for Multi-Instance Mode 421

CHAPTER 10 **High Availability** 423
About Secure Firewall Threat Defense High Availability 423
High Availability System Requirements 423
Hardware Requirements 424
Software Requirements 424
License Requirements for Firewall Threat Defense Devices in a High Availability Pair 424
Failover and Stateful Failover Links 425
Failover Link 425
Stateful Failover Link 427
Avoiding Interrupted Failover and Data Links 427
MAC Addresses and IP Addresses in High Availability 428
Stateful Failover 430
Supported Features 430
Unsupported Features 431
Bridge Group Requirements for High Availability 432
Failover Health Monitoring 432
Unit Health Monitoring 432
Heartbeat Module Redundancy 433

Interface Monitoring 433
Failover Triggers and Detection Timing 435
About Active/Standby Failover 436
Primary/Secondary Roles and Active/Standby Status 436
Active Unit Determination at Startup 436
Failover Events 436
Config-Sync Optimization 437
Requirements and Prerequisites for High Availability 438
Guidelines for High Availability 439
Add a High Availability Pair 441
Configure Optional High Availability Parameters 444
Configure Standby IP Addresses and Interface Monitoring 444
Edit High Availability Failover Criteria 444
Configure Virtual MAC Addresses 445
Manage High Availability 446
Switch the Active Peer in the Firewall Threat Defense High Availability Pair 446
Refresh Node Status for a Single Firewall Threat Defense High Availability Pair 446
Suspend and Resume High Availability 447
Replace a Unit in Firewall Threat Defense High Availability Pair 448
Replace a Primary Firewall Threat Defense HA Unit with no Backup 448
Replace a Secondary Firewall Threat Defense HA Unit with no Backup 449
Break a High Availability Pair 450
Unregister a High Availability Pair and Register to a New Firewall Management Center 451
Monitoring High Availability 453
View Failover History 453
View Stateful Failover Statistics 453
Troubleshoot Configuration Sync Failure 453
History for High Availability 454

CHAPTER 11 **Clustering for the Secure Firewall 3100/4200** 457
About Clustering for the Secure Firewall 3100/4200 457
How the Cluster Fits into Your Network 457
Control and Data Node Roles 458
Cluster Interfaces 458

Cluster Control Link 458
Configuration Replication 458
Management Network 458
Licenses for Clustering 459
Requirements and Prerequisites for Clustering 459
Guidelines for Clustering 460
Configure Clustering 464
About Cluster Interfaces 464
Cluster Control Link 464
Spanned EtherChannels (Recommended) 466
Individual Interfaces (Routed Firewall Mode Only) 469
Cable and Add Devices to the Firewall Management Center 471
Create a Cluster 472
Configure Interfaces 481
Configure Spanned EtherChannels 481
Configure Individual Interfaces 482
Configure Cluster Health Monitor Settings 484
Manage Cluster Nodes 489
Add a New Cluster Node 489
Break a Node 491
Break the Cluster 492
Disable Clustering 493
Rejoin the Cluster 494
Change the Control Node 494
Edit the Cluster Configuration 495
Reconcile Cluster Nodes 497
Unregister the Cluster or Nodes and Register to a New Firewall Management Center 498
Monitoring the Cluster 499
Cluster Health Monitor Dashboard 502
Viewing Cluster Health 503
Cluster Metrics 504
Troubleshooting the Cluster 505
Perform a Ping on the Cluster Control Link 505
Examples for Clustering 507

Firewall on a Stick 508
Traffic Segregation 509
Reference for Clustering 509
Firewall Threat Defense Features and Clustering 509
Unsupported Features with Clustering 509
Centralized Features for Clustering 510
Connection Settings and Clustering 511
FTP and Clustering 511
Multicast Routing in Individual Interface Mode 511
Multicast Routing in Individual Interface Mode 511
NAT and Clustering 511
Dynamic Routing 513
Dynamic Routing in Individual Interface Mode 514
SIP Inspection and Clustering 515
SNMP and Clustering 515
Syslog and Clustering 515
Cisco TrustSec and Clustering 515
VPN and Clustering 515
Performance Scaling Factor 515
Control Node Election 516
High Availability Within the Cluster 516
Node Health Monitoring 516
Interface Monitoring 517
Status After Failure 517
Rejoining the Cluster 517
Data Path Connection State Replication 518
How the Cluster Manages Connections 518
Connection Roles 519
New Connection Ownership 520
Sample Data Flow for TCP 520
Sample Data Flow for ICMP and UDP 521
History for Clustering 522

CHAPTER 12 **Clustering for Threat Defense Virtual in a Private Cloud** 525

About Threat Defense Virtual Clustering in the Private Cloud 525
How the Cluster Fits into Your Network 525
Control and Data Node Roles 526
Individual Interfaces 526
Policy-Based Routing 527
Equal-Cost Multi-Path Routing 527
Cluster Control Link 528
Cluster Control Link Traffic Overview 528
Configuration Replication 529
Management Network 529
Licenses for Threat Defense Virtual Clustering 529
Requirements and Prerequisites for Threat Defense Virtual Clustering 530
Guidelines for Threat Defense Virtual Clustering 531
Configure Threat Defense Virtual Clustering 532
Add Nodes to the Management Center 532
Create a Cluster 532
Configure Interfaces 541
Configure Cluster Health Monitor Settings 541
Manage Cluster Nodes 546
Add a New Cluster Node 546
Break a Node 548
Break the Cluster 549
Disable Clustering 550
Rejoin the Cluster 551
Change the Control Node 551
Edit the Cluster Configuration 552
Reconcile Cluster Nodes 554
Delete (Unregister) the Cluster or Nodes and Register to a New Firewall Management Center 555
Monitoring the Cluster 556
Cluster Health Monitor Dashboard 559
Viewing Cluster Health 560
Cluster Metrics 561
Troubleshooting the Cluster 562
Perform a Ping on the Cluster Control Link 562

Reference for Clustering 564
Threat Defense Features and Clustering 564
Unsupported Features and Clustering 564
Centralized Features for Clustering 565
Connection Settings and Clustering 566
Dynamic Routing and Clustering 566
FTP and Clustering 566
NAT and Clustering 567
SIP Inspection and Clustering 568
SNMP and Clustering 568
Syslog and Clustering 568
Cisco Trustsec and Clustering 569
VPN and Clustering 569
Performance Scaling Factor 569
Control Node Election 569
High Availability within the Cluster 570
Node Health Monitoring 570
Interface Monitoring 570
Status After Failure 570
Rejoining the Cluster 571
Data Path Connection State Replication 571
How the Cluster Manages Connections 572
Connection Roles 572
New Connection Ownership 573
Sample Data Flow for TCP 573
Sample Data Flow for ICMP and UDP 574
History for Threat Defense Virtual Clustering in a Private Cloud 576

CHAPTER 13 **Clustering for Threat Defense Virtual in a Public Cloud 579**
About Threat Defense Virtual Clustering in the Public Cloud 580
How the Cluster Fits into Your Network 580
Individual Interfaces 580
Control and Data Node Roles 581
Cluster Control Link 581

Cluster Control Link Traffic Overview 581
Configuration Replication 582
Management Network 582
Licenses for Threat Defense Virtual Clustering 582
Requirements and Prerequisites for Threat Defense Virtual Clustering 583
Guidelines for Threat Defense Virtual Clustering 584
Deploy the Cluster in AWS 586
AWS Gateway Load Balancer and Geneve Single-Arm Proxy 586
Sample Topologies 587
End-to-End Process for Deploying Threat Defense Virtual Cluster on AWS 589
Templates 591
Deploy the Stack in AWS Using a CloudFormation Template 592
Autoscale Parameter Configuration 604
Configure IMDSv2 Required Mode in Firewall Threat Defense Virtual Clustering by Updating Stack 604
Deploy the Cluster in AWS Manually 605
Create the Day0 Configuration for AWS 605
Deploy Cluster Nodes 607
Configure Target Failover for Secure Firewall Threat Defense Virtual Clustering with GWLB in AWS 608
Enable Target Failover for Secure Firewall Threat Defense Virtual Clustering in AWS 609
Deploy the Cluster in Azure 610
Sample Topology for GWLB-based Cluster Deployment 610
Azure Gateway Load Balancer and Paired Proxy 611
End-to-End Process for Deploying Threat Defense Virtual Cluster in Azure with GWLB 612
Templates 614
Prerequisites 614
Deploy Cluster on Azure with GWLB Using an Azure Resource Manager Template 615
Sample Topology for NLB-based Cluster Deployment 618
End-to-End Process for Deploying Threat Defense Virtual Cluster in Azure with NLB 618
Templates 620
Prerequisites 621
Deploy Cluster on Azure with NLB Using an Azure Resource Manager Template 621
Deploy the Cluster in Azure Manually 623

Create the Day0 Configuration for Azure 623
Deploy Cluster Nodes Manually - GWLB-based Deployment 625
Deploy Cluster Nodes Manually - NLB-based Deployment 627
Troubleshooting Cluster Deployment in Azure 628
Firewall Threat Defense Virtual Clustering Autoscale Solution on Azure 629
Sample Topologies 630
Prerequisites 632
Autoscale Logic for Firewall Threat Defense Virtual Clustering in Azure 632
Deployment and Infrastructure Templates on GitHub 634
Input Parameters 634
Firewall Threat Defense Virtual Cluster with Autoscale Deployment Process and Resources 643
Deploy the Firewall Threat Defense Virtual Cluster with Autoscale Solution 644
Deploy Azure Functions App 650
Update the Azure Logic App 650
Deploy the Cluster in GCP 653
Sample Topology 654
End-to-End Process for Deploying Threat Defense Virtual Cluster in GCP 654
Templates 656
Deploy the Instance Group in GCP Using an Instance Template 656
Deploy the Cluster in GCP Manually 658
Create the Day0 Configuration for GCP 658
Deploy Cluster Nodes Manually 660
Allow Health Checks for GCP Network Load Balancers 660
Add the Cluster to the Management Center (Manual Deployment) 662
Configure Cluster Health Monitor Settings 669
Manage Cluster Nodes 674
Disable Clustering 674
Rejoin the Cluster 674
Reconcile Cluster Nodes 674
Unregister the Cluster or Nodes and Register to a New Firewall Management Center 675
Monitoring the Cluster 676
Cluster Health Monitor Dashboard 679
Viewing Cluster Health 680
Cluster Metrics 681

Troubleshooting the Cluster 682
Perform a Ping on the Cluster Control Link 682
Upgrading the Cluster 684
Reference for Clustering 685
Threat Defense Features and Clustering 685
Unsupported Features and Clustering 685
Centralized Features for Clustering 686
Cisco Trustsec and Clustering 686
Connection Settings and Clustering 686
Dynamic Routing and Clustering 687
FTP and Clustering 687
NAT and Clustering 688
SIP Inspection and Clustering 689
SNMP and Clustering 689
Syslog and Clustering 689
VPN and Clustering 689
Performance Scaling Factor 690
Control Node Election 690
High Availability within the Cluster 691
Node Health Monitoring 691
Interface Monitoring 691
Status After Failure 691
Rejoining the Cluster 691
Data Path Connection State Replication 692
How the Cluster Manages Connections 692
Connection Roles 692
New Connection Ownership 694
Sample Data Flow for TCP 694
Sample Data Flow for ICMP and UDP 695
History for Threat Defense Virtual Clustering in the Public Cloud 696

CHAPTER 14 **Clustering for the Firepower 4100/9300** 699
About Clustering on the Firepower 4100/9300 Chassis 699
Bootstrap Configuration 700

Cluster Members 700
Cluster Control Link 700
Size the Cluster Control Link 701
Cluster Control Link Redundancy 701
Cluster Control Link Reliability for Inter-Chassis Clustering 702
Cluster Control Link Network 702
Management Network 702
Management Interface 702
Cluster Interfaces 702
Spanned EtherChannels 702
Connecting to a Redundant Switch System 703
Configuration Replication 703
Licenses for Clustering 703
Requirements and Prerequisites for Clustering 704
Clustering Guidelines and Limitations 707
Configure Clustering 711
FXOS: Add a Firewall Threat Defense Cluster 711
Create a Firewall Threat Defense Cluster 711
Add More Cluster Nodes 722
Firewall Management Center: Add a Cluster 725
Firewall Management Center: Configure Cluster, Data Interfaces 731
Firewall Management Center: Configure Cluster Health Monitor Settings 733
FXOS: Remove a Cluster Node 737
Firewall Management Center: Manage Cluster Members 739
Add a New Cluster Member 739
Replace a Cluster Member 739
Deactivate a Member 741
Rejoin the Cluster 741
Unregister a Data Node 742
Change the Control Unit 743
Reconcile Cluster Members 744
Firewall Management Center: Monitoring the Cluster 744
Cluster Health Monitor Dashboard 746
Viewing Cluster Health 747

Cluster Metrics 748
Firewall Management Center: Troubleshooting the Cluster 749
Perform a Ping on the Cluster Control Link 750
Examples for Clustering 751
Firewall on a Stick 752
Traffic Segregation 753
Reference for Clustering 753
Firewall Threat Defense Features and Clustering 753
Unsupported Features with Clustering 753
Centralized Features for Clustering 754
Connection Settings 755
Dynamic Routing and Clustering 755
FTP and Clustering 756
Multicast Routing and Clustering 756
NAT and Clustering 756
SIP Inspection and Clustering 757
SNMP and Clustering 757
Syslog and Clustering 758
TLS/SSL Connections and Clustering 758
Cisco TrustSec and Clustering 758
VPN and Clustering 758
Performance Scaling Factor 758
Control Unit Election 758
High Availability Within the Cluster 759
Chassis-Application Monitoring 759
Unit Health Monitoring 759
Interface Monitoring 760
Decorator Application Monitoring 760
Status After Failure 760
Rejoining the Cluster 760
Data Path Connection State Replication 761
How the Cluster Manages Connections 762
Connection Roles 762
New Connection Ownership 763

Sample Data Flow for TCP 763
Sample Data Flow for ICMP and UDP 764
History for Clustering 765

PART III **Interfaces and Device Settings** 771

CHAPTER 15 **Interface Overview** 773
Management Interface 773
Management Interface 773
Diagnostic Interface 774
Interface Mode and Types 774
Security Zones and Interface Groups 775
Auto-MDI/MDIX Feature 776
Redundant Interfaces (Deprecated) 777
Default Settings for Interfaces 777
Create Security Zone and Interface Group Objects 777
Enable the Physical Interface and Configure Ethernet Settings 778
Configure EtherChannel Interfaces 781
About EtherChannels 781
About EtherChannels 781
Guidelines for EtherChannels 784
Configure an EtherChannel 786
Sync Interface Changes with the Firewall Management Center 789
Manage the Network Module for the Secure Firewall 3100/4200 790
Configure Breakout Ports 791
Add a Network Module 794
Hot Swap the Network Module 797
Replace the Network Module with a Different Type 799
Remove the Network Module 802
Merge the Management and Diagnostic Interfaces 805
Unmerge the Management Interface 811
History for Interfaces 812

CHAPTER 16 **Regular Firewall Interfaces** 817

Requirements and Prerequisites for Regular Firewall Interfaces 817
Configure Firepower 1010 and Secure Firewall 1210/1220 Switch Ports 818
About Switch Ports 818
Understanding Switch Ports and Interfaces 818
Auto-MDI/MDIX Feature 819
Guidelines and Limitations for Switch Ports 819
Configure Switch Ports and Power Over Ethernet 820
Enable or Disable Switch Port Mode 821
Configure a VLAN Interface 821
Configure Switch Ports as Access Ports 823
Configure Switch Ports as Trunk Ports 825
Configure Power Over Ethernet 827
Configure Loopback Interfaces 828
About Loopback Interfaces 829
Guidelines and Limitations for Loopback Interfaces 829
Configure a Loopback Interface 830
Rate-Limit Traffic to the Loopback Interface 830
Configure VLAN Subinterfaces and 802.1Q Trunking 834
Guidelines and Limitations for VLAN Subinterfaces 834
Maximum Number of VLAN Subinterfaces by Device Model 835
Add a Subinterface 836
Configure VXLAN Interfaces 838
About VXLAN Interfaces 838
Encapsulation 838
VXLAN Tunnel Endpoint 838
VTEP Source Interface 839
VNI Interfaces 839
VXLAN Packet Processing 840
Peer VTEPs 840
VXLAN Use Cases 841
Requirements and Prerequisites for VXLAN Interfaces 846
Guidelines for VXLAN Interfaces 847
Configure VXLAN or Geneve Interfaces 848
Configure VXLAN Interfaces 848

Configure Geneve Interfaces 850
Allow Gateway Load Balancer Health Checks 852
Configure Routed and Transparent Mode Interfaces 853
About Routed and Transparent Mode Interfaces 853
Dual IP Stack (IPv4 and IPv6) 853
31-Bit Subnet Mask 853
Guidelines and Limitations for Routed and Transparent Mode Interfaces 854
Configure Routed Mode Interfaces 856
Configure Bridge Group Interfaces 861
Configure General Bridge Group Member Interface Parameters 861
Configure the Bridge Virtual Interface (BVI) 863
Configure IPv6 Addressing 864
About IPv6 864
Configure the IPv6 Prefix Delegation Client 865
Configure a Global IPv6 Address 868
Configure IPv6 Neighbor Discovery 873
Configure Advanced Interface Settings 875
About Advanced Interface Configuration 875
About MAC Addresses 875
About the MTU 877
About the TCP MSS 878
ARP Inspection for Bridge Group Traffic 879
MAC Address Table 879
Default Settings 880
Guidelines for ARP Inspection and the MAC Address Table 880
Configure the MTU 880
Configure the MAC Address 881
Add a Static ARP Entry 882
Add a Static MAC Address and Disable MAC Learning for a Bridge Group 883
Set Security Configuration Parameters 884
History for Regular Firewall Interfaces 886

CHAPTER 17 **Inline Sets and Passive Interfaces** 893
About IPS Interfaces 893

Inline Sets 893
Multiple Inline Pairs and Asynchronous Routing 894
Passive Interfaces 895
About Hardware Bypass for Inline Sets 896
Hardware Bypass Triggers 896
Hardware Bypass Switchover 896
Snort Fail Open vs. Hardware Bypass 896
Hardware Bypass Status 897
Requirements and Prerequisites for Inline Sets 897
Guidelines for Inline Sets and Passive Interfaces 898
Configure a Passive Interface 900
Configure an Inline Set 901
History for Inline Sets and Passive Interfaces 905

CHAPTER 18 DHCP and DDNS 909
About DHCP and DDNS Services 909
About the DHCPv4 Server 909
DHCP Options 909
About the DHCPv6 Stateless Server 910
About the DHCP Relay Agent 910
Requirements and Prerequisites for DHCP and DDNS 911
Guidelines for DHCP and DDNS Services 911
Configure the DHCPv4 Server 913
Configure the DHCPv6 Stateless Server 914
Create the DHCP IPv6 Pool 914
Enable the DHCPv6 Stateless Server 917
Configure the DHCP Relay Agent 919
Configure Dynamic DNS 920
History for DHCP and DDNS 926

CHAPTER 19 SNMP for the Firepower 1000 927
About SNMP for the Firepower 1000 927
Enabling SNMP and Configuring SNMP Properties for Firepower 1000 927
Creating an SNMP Trap for Firepower 1000 929

Creating an SNMP User for Firepower 1000 930

CHAPTER 20 **Quality of Service** 933

Introduction to QoS 933

About QoS Policies 933

Requirements and Prerequisites for QoS 934

Rate Limiting with QoS Policies 934

Creating a QoS Policy 935

Setting Target Devices for a QoS Policy 936

Configuring QoS Rules 936

QoS Rule Components 937

QoS Rule Conditions 938

Interface Rule Conditions 938

Network Rule Conditions 939

User Rule Conditions 939

Application Rule Conditions 939

Port Rule Conditions 940

URL Rule Conditions 942

Custom SGT Rule Conditions 942

ISE SGT vs Custom SGT Rule Conditions 942

Autotransition from Custom SGTs to ISE SGTs 943

History for QoS 943

CHAPTER 21 **Platform Settings** 945

Introduction to Platform Settings 946

Requirements and Prerequisites for Platform Settings Policies 946

Manage Platform Settings Policies 946

Chassis Platform Settings 947

ARP Inspection 947

Banner 948

DNS 949

External Authentication 952

Enable Virtual-Router-Aware Interface for External Authentication of Platform 957

Fragment Settings 958

HTTP Access 959
ICMP Access 960
NetFlow 962
Add Collector in NetFlow 963
Add Traffic Class to NetFlow 963
SSH Access 964
SMTP Server 966
SNMP 966
About SNMP 968
SNMP Terminology 968
MIBs and Traps 969
Supported Tables and Objects in MIBs 969
Add SNMPv3 Users 973
Add SNMP Hosts 976
Configure SNMP Traps 978
Configure SSL Settings 980
About SSL Settings 980
Syslog 984
About Syslog 984
Severity Levels 985
Syslog Message Filtering 986
Syslog Message Classes 986
Guidelines for Logging 989
Configure Syslog Logging for Firewall Threat Defense Devices 990
Firewall Threat Defense Platform Settings That Apply to Security Event Syslog Messages 991
Enable Logging and Configure Basic Settings 991
Enable Logging Destinations 993
Send Syslog Messages to an E-mail Address 994
Create a Custom Event List 995
Limit the Rate of Syslog Message Generation 996
Configure Syslog Settings 997
Configure a Syslog Server 998
Timeouts 1000
Time Synchronization 1002

Time Zone 1003
UCAPL/CC Compliance 1004
Performance Profile 1005
History for Platform Settings 1006

CHAPTER 22 **Network Address Translation** 1015
Why Use NAT? 1015
NAT Basics 1016
NAT Terminology 1016
NAT Types 1016
NAT in Routed and Transparent Mode 1017
NAT in Routed Mode 1017
NAT in Transparent Mode or Within a Bridge Group 1018
Auto NAT and Manual NAT 1019
Auto NAT 1019
Manual NAT 1019
Comparing Auto NAT and Manual NAT 1020
NAT Rule Order 1020
NAT Interfaces 1022
NAT Exemption 1023
Configuring Routing for NAT 1023
Addresses on the Same Network as the Mapped Interface 1024
Addresses on a Unique Network 1024
The Same Address as the Real Address (Identity NAT) 1024
Requirements and Prerequisites for NAT Policies 1025
Guidelines for NAT 1025
Firewall Mode Guidelines for NAT 1025
IPv6 NAT Guidelines 1025
IPv6 NAT Best Practices 1026
NAT Support for Inspected Protocols 1026
FQDN Destination Guidelines 1028
Additional Guidelines for NAT 1029
Manage NAT Policies 1031
Creating NAT Policies 1032

Configuring NAT Policy Targets 1032
Configure NAT for Threat Defense 1033
Customizing NAT Rules for Multiple Devices 1034
Searching and Filtering the NAT Rule Table 1036
Enabling, Disabling, or Deleting Multiple Rules 1037
Dynamic NAT 1038
About Dynamic NAT 1038
Dynamic NAT Disadvantages and Advantages 1039
Configure Dynamic Auto NAT 1040
Configure Dynamic Manual NAT 1041
Dynamic PAT 1043
About Dynamic PAT 1043
Dynamic PAT Disadvantages and Advantages 1044
PAT Pool Object Guidelines 1044
Configure Dynamic Auto PAT 1045
Configure Dynamic Manual PAT 1048
Configure PAT with Port Block Allocation 1051
Static NAT 1053
About Static NAT 1053
Configure Static Auto NAT 1057
Configure Static Manual NAT 1058
Identity NAT 1061
Configure Identity Auto NAT 1062
Configure Identity Manual NAT 1063
NAT Rule Properties for Firewall Threat Defense 1066
Interface Objects NAT Properties 1066
Translation Properties for Auto NAT 1067
Translation Properties for Manual NAT 1068
PAT Pool NAT Properties 1069
Advanced NAT Properties 1070
Translating IPv6 Networks 1071
NAT64/46: Translating IPv6 Addresses to IPv4 1072
NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet 1072
NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet and DNS Translation 1074

NAT66: Translating IPv6 Addresses to Different IPv6 Addresses 1078
NAT66 Example, Static Translation between Networks 1078
NAT66 Example, Simple IPv6 Interface PAT 1081
Monitoring NAT 1084
Examples for NAT 1085
Providing Access to an Inside Web Server (Static Auto NAT) 1085
Dynamic Auto NAT for Inside Hosts and Static NAT for an Outside Web Server 1088
Inside Load Balancer with Multiple Mapped Addresses (Static Auto NAT, One-to-Many) 1092
Single Address for FTP, HTTP, and SMTP (Static Auto NAT-with-Port-Translation) 1095
Different Translation Depending on the Destination (Dynamic Manual PAT) 1101
Different Translation Depending on the Destination Address and Port (Dynamic Manual PAT) 1107
NAT and Site-to-Site VPN 1113
Rewriting DNS Queries and Responses Using NAT 1117
DNS64 Reply Modification 1118
DNS Reply Modification, DNS Server on Outside 1126
DNS Reply Modification, DNS Server on Host Network 1129
History for Firewall Threat Defense NAT 1133

CHAPTER 23 **Alarms for the Cisco ISA 3000** 1135
About Alarms 1135
Alarm Input Interfaces 1136
Alarm Output Interface 1136
Syslog Alarms 1137
SNMP Alarms 1137
Defaults for Alarms 1137
Requirements and Prerequisites for Alarms 1138
Configure the Alarms for the ISA 3000 1138
Configure Alarm Input Contacts 1138
Configure Power Supply Alarms 1141
Configure Temperature Alarms 1143
Monitoring Alarms 1146
Monitoring Alarm Status 1146
Monitoring Syslog Messages for Alarms 1146
Turning Off the External Alarm 1147

History for Alarms 1147

PART IV **Routing** 1149

CHAPTER 24 **Static and Default Routes** 1151

About Static and Default Routes 1151

Default Route 1151

Static Routes 1152

Route to null0 Interface to Drop Unwanted Traffic 1152

Route Priorities 1152

Transparent Firewall Mode and Bridge Group Routes 1152

Static Route Tracking 1153

Requirements and Prerequisites for Static Routes 1153

Guidelines for Static and Default Routes 1154

Add a Static Route 1154

Reference for Routing 1156

Path Determination 1156

Supported Route Types 1156

Static Versus Dynamic 1156

Single-Path Versus Multipath 1157

Flat Versus Hierarchical 1157

Link-State Versus Distance Vector 1157

Supported Internet Protocols for Routing 1157

Routing Table 1158

How the Routing Table Is Populated 1158

How Forwarding Decisions Are Made 1160

Dynamic Routing and High Availability 1161

Dynamic Routing in Clustering 1161

Dynamic Routing in Individual Interface Mode 1162

Routing Table for Management Traffic 1163

Equal-Cost Multi-Path (ECMP) Routing 1164

About Route Maps 1165

Permit and Deny Clauses 1165

Match and Set Clause Values 1165

CHAPTER 25 **Virtual Routers** 1167

About Virtual Routers and Virtual Routing and Forwarding (VRF) 1167

About Virtual Routers and Dynamic VTI 1168

How to Configure a Virtual Router with Dynamic VTI 1168

Applications of Virtual Routers 1168

Global and User-Defined Virtual Routers 1169

Configuring Policies to be Virtual-Router-Aware 1170

Interconnecting Virtual Routers 1171

Overlapping IP Addresses 1172

Configuring SNMP on User-Defined Virtual Routers 1173

Maximum Number of Virtual Routers By Device Model 1174

Requirements and Prerequisites for Virtual Routers 1175

Guidelines and Limitations for Virtual Routers 1175

Modifications to the Firewall Management Center Web Interface - Routing Page 1177

Manage Virtual Routers 1178

Create a Virtual Router 1178

Configure a Virtual Router 1179

Modify a Virtual Router 1180

Remove Virtual Routers 1181

Monitoring Virtual Routers 1182

Configuration Examples for Virtual Routers 1182

How to Route to a Distant Server through Virtual Routers 1182

How to Provide Internet Access with Overlapping Address Spaces 1187

How to Allow RA VPN Access to Internal Networks in Virtual Routing 1194

How to Secure Traffic from Networks in Multiple Virtual Routers over a Site-to-Site VPN 1197

How to Secure Traffic from Networks with Multiple Virtual Routers over a Site-to-Site VPN with Dynamic VTI 1201

How to Route Traffic between Two Overlapping Network Host in Virtual Routing 1203

How to Manage Overlapping Segments in Routed Firewall Mode with BVI Interfaces 1206

How to Configure User Authentication with Overlapping Networks 1210

How to Interconnect Virtual Routers using BGP 1216

History for Virtual Routers 1221

CHAPTER 26 ECMP 1223

About ECMP 1223

Guidelines and Limitations for ECMP 1223

Manage ECMP Page 1224

Create an ECMP Zone 1225

Configure an Equal Cost Static Route 1226

Modify an ECMP Zone 1227

Remove an ECMP Zone 1227

Configuration Example for ECMP 1228

History for ECMP in Secure Firewall Threat Defense 1231

CHAPTER 27 Bidirectional Forwarding Detection Routing 1233

About BFD Routing 1233

Guidelines for BFD Routing 1233

Configure BFD 1235

Configure BFD Policies 1235

Configure Single-Hop BFD Policies 1236

Configure Multi-Hop BFD Policies 1236

History for BFD Routing 1237

CHAPTER 28 OSPF 1239

OSPF 1239

About OSPF 1239

OSPF Support for Fast Hello Packets 1241

Prerequisites for OSPF Support for Fast Hello Packets 1241

OSPF Hello Interval and Dead Interval 1241

OSPF Fast Hello Packets 1241

Benefits of OSPF Fast Hello Packets 1241

Implementation Differences Between OSPFv2 and OSPFv3 1242

Requirements and Prerequisites for OSPF 1242

Guidelines for OSPF 1242

Configure OSPFv2 1245

Configure OSPF Areas, Ranges, and Virtual Links 1245

Configure OSPF Redistribution 1248
Configure OSPF Inter-Area Filtering 1249
Configure OSPF Filter Rules 1250
Configure OSPF Summary Addresses 1251
Configure OSPF Interfaces and Neighbors 1252
Configure OSPF Advanced Properties 1254
Configure OSPFv3 1257
Configure OSPFv3 Areas, Route Summaries, and Virtual Links 1257
Configure OSPFv3 Redistribution 1259
Configure OSPFv3 Summary Prefixes 1260
Configure OSPFv3 Interfaces, Authentication, and Neighbors 1261
Configure OSPFv3 Advanced Properties 1264
History for OSPF 1266

CHAPTER 29 EIGRP 1267
About EIGRP Routing 1267
Requirements and Prerequisites for EIGRP 1268
Guidelines and Limitations of EIGRP Routing 1268
Configure EIGRP 1270
Configure EIGRP Settings 1270
Configure EIGRP Neighbors Settings 1271
Configure EIGRP Filter Rules Settings 1271
Configure EIGRP Redistribution Settings 1272
Configure EIGRP Summary Address Settings 1273
Configure EIGRP Interfaces Settings 1274
Configure EIGRP Advanced Settings 1274
History for EIGRP 1276

CHAPTER 30 BGP 1277
About BGP 1277
Routing Table Changes 1277
When to Use BGP 1278
BGP Path Selection 1279
BGP Multipath 1279

Requirements and Prerequisites for BGP 1280
Guidelines for BGP 1280
Configure BGP 1281
Configure BGP Basic Settings 1282
Configure BGP General Settings 1284
Configure BGP Neighbor Settings 1285
Configure BGP Aggregate Address Settings 1289
Configure BGPv4 Filtering Settings 1290
Configure BGP Network Settings 1291
Configure BGP Redistribution Settings 1291
Configure BGP Route Injection Settings 1292
Configure BGP Route Import/Export Settings 1293
History for BGP 1295

CHAPTER 31 RIP 1297
About RIP 1297
Routing Update Process 1297
RIP Routing Metric 1298
RIP Stability Features 1298
RIP Timers 1298
Requirements and Prerequisites for RIP 1299
Guidelines for RIP 1299
Configure RIP 1300

CHAPTER 32 Multicast 1303
About Multicast Routing 1303
IGMP Protocol 1303
Stub Multicast Routing 1304
PIM Multicast Routing 1304
PIM Source Specific Multicast Support 1305
Multicast Bidirectional PIM 1305
PIM Bootstrap Router (BSR) 1306
PIM Bootstrap Router (BSR) Terminology 1306
Multicast Group Concept 1307

Multicast Addresses 1307
Clustering 1307
Requirements and Prerequisites for Multicast Routing 1307
Guidelines for Multicast Routing 1307
Configure IGMP Features 1308
Enable Multicast Routing 1309
Configure IGMP Protocol 1309
Configure IGMP Access Groups 1311
Configure IGMP Static Groups 1311
Configure IGMP Join Groups 1312
Configure PIM Features 1313
Configure PIM Protocol 1313
Configure PIM Neighbor Filters 1314
Configure PIM Bidirectional Neighbor Filters 1315
Configure PIM Rendezvous Points 1316
Configure PIM Route Trees 1317
Configure PIM Request Filters 1317
Configure the Secure Firewall Threat Defense Device as a Candidate Bootstrap Router 1318
Configure Multicast Routes 1319
Configure Multicast Boundary Filters 1320

CHAPTER 33 **Policy Based Routing** 1323
About Policy Based Routing 1323
Guidelines and Limitations for Policy Based Routing 1325
Path Monitoring 1326
Configure Path Monitoring Settings 1328
Configure Policy-Based Routing Policy 1329
Add Path Monitoring Dashboard 1332
Create Custom Application Detector for PBR 1332
Configuration Example for PBR with Custom Application Detector 1333
Configuration Example for Policy Based Routing 1339
Configuration Example for PBR with Path Monitoring 1344
History for Policy Based Routing 1346

PART V **Objects and Certificates** 1349

CHAPTER 34 **Object Management** 1351

Introduction to Objects 1352

The Object Manager 1354

Importing Objects 1354

Editing Objects 1357

Viewing Objects and Their Usage 1358

Filtering Objects or Object Groups 1359

Object Groups 1359

Grouping Reusable Objects 1360

Object Overrides 1361

Managing Object Overrides 1362

Allowing Object Overrides 1362

Adding Object Overrides 1362

Editing Object Overrides 1363

AAA Server 1364

Add a RADIUS Server Group 1364

RADIUS Server Group Options 1364

RADIUS Server Options 1366

RADIUS Server-Enabled Message Authenticator Compatibility Matrix 1367

Add a Single Sign-on Server 1367

Access List 1369

Configure Extended ACL Objects 1370

Configure a Service Access Object 1372

Configure Standard ACL Objects 1374

Address Pools 1374

Application Filters 1376

AS Path 1376

BFD Template 1377

Cipher Suite List 1378

Creating Cipher Suite Lists 1378

Community List 1379

Extended Community 1380
DHCP IPv6 Pool 1381
Distinguished Name 1382
Creating Distinguished Name Objects 1384
DNS Server Group 1384
Creating DNS Server Group Objects 1384
External Attributes 1385
Dynamic Objects 1385
Create Dynamic Objects for the First Time 1386
Create Dynamic Objects with the Embedded Cisco Secure Dynamic Attributes Connector 1387
Create Dynamic Objects with Cisco Security Cloud Control 1388
Create Dynamic Objects with the On-Premises Cisco Secure Dynamic Attributes Connector 1389
Work With Dynamic Objects 1390
Dynamic Object Mappings 1390
About API-Created Dynamic Objects 1390
Security Group Tag 1392
Creating Security Group Tag Objects 1392
File List 1392
Source Files for File Lists 1393
Adding Individual SHA-256 Values to File Lists 1394
Uploading Individual Files to File Lists 1394
Uploading Source Files to File Lists 1395
Editing SHA-256 Values in File Lists 1396
Downloading Source Files from File Lists 1397
FlexConfig 1397
Geolocation 1398
Creating Geolocation Objects 1398
Interface 1398
Key Chain 1399
Creating Key Chain Objects 1399
Network 1401
Network Wildcard Mask 1402
Creating Network Objects 1403
Importing Network Objects 1404

PKI 1404
Internal Certificate Authority Objects 1405
CA Certificate and Private Key Import 1405
Importing a CA Certificate and Private Key 1406
Generating a New CA Certificate and Private Key 1406
New Signed Certificates 1407
Creating an Unsigned CA Certificate and CSR 1407
Uploading a Signed Certificate Issued in Response to a CSR 1407
CA Certificate and Private Key Downloads 1408
Downloading a CA Certificate and Private Key 1408
Trusted Certificate Authority Objects 1409
Trusted CA Object 1409
Adding a Trusted CA Object 1409
Certificate Revocation Lists in Trusted CA Objects 1410
Adding a Certificate Revocation List to a Trusted CA Object 1410
External Certificate Objects 1411
Adding External Certificate Objects 1411
Internal Certificate Objects 1412
Adding Internal Certificate Objects 1412
Certificate Enrollment Objects 1413
Adding Certificate Enrollment Objects 1414
Add Certificate Enrollment 1416
Certificate Enrollment Object EST Options 1416
Certificate Enrollment Object SCEP Options 1417
Certificate Enrollment Object Certificate Parameters 1418
Certificate Enrollment Object Key Options 1419
Certificate Enrollment Object Revocation Options 1420
Policy List 1421
Port 1422
Creating Port Objects 1423
Importing Port Objects 1424
Prefix List 1424
Configure IPv6 Prefix List 1424
Configure IPv4 Prefix List 1425

Route Map 1426
Security Intelligence 1429
How to Modify Security Intelligence Objects 1431
Global and Domain Security Intelligence Lists 1431
Security Intelligence Lists and Multitenancy 1431
Add Entries to Global Security Intelligence Lists 1432
Delete Entries from Global Security Intelligence Lists 1433
List and Feed Updates for Security Intelligence 1434
Changing the Update Frequency for Security Intelligence Feeds 1434
Custom Security Intelligence Lists and Feeds 1435
Custom Lists and Feeds: Requirements 1435
URL Lists and Feeds: URL Syntax and Matching Criteria 1435
Custom Security Intelligence Feeds 1437
Custom Security Intelligence Lists 1438
Sinkhole 1440
Creating Sinkhole Objects 1440
SLA Monitor 1441
Time Range 1442
Creating Time Range Objects 1442
Time Zone 1444
Tunnel Zone 1444
URL 1444
Creating URL Objects 1445
Variable Set 1446
Variable Sets in Intrusion Policies 1447
Variables 1447
Predefined Default Variables 1448
Network Variables 1450
Port Variables 1452
Advanced Variables 1453
Variable Reset 1453
Adding Variables to Sets 1454
Nesting Variables 1455
Managing Variable Sets 1457

Creating Variable Sets 1457
Managing Variables 1458
Adding Variables 1459
Editing Variables 1460
VLAN Tag 1461
Creating VLAN Tag Objects 1461
VPN 1462
Certificate Map Objects 1462
Secure Client Custom Attributes Objects 1463
Add Secure Client Custom Attributes Objects 1463
Add Custom Attributes to a Group Policy 1465
Firewall Threat Defense Group Policy Objects 1465
Configure Group Policy Objects 1466
Group Policy General Options 1466
Group Policy Secure Client Options 1469
Group Policy Advanced Options 1472
Firewall Threat Defense IPsec Proposals 1473
Configure IKEv1 IPsec Proposal Objects 1473
Configure IKEv2 IPsec Proposal Objects 1474
Firewall Threat Defense IKE Policies 1475
Configure IKEv1 Policy Objects 1475
Configure IKEv2 Policy Objects 1476
Secure Client Customization 1478
File Objects 1478
History for Object Management 1480

CHAPTER 35 **Certificates** 1485
Requirements and Prerequisites for Certificates 1485
Secure Firewall Threat Defense VPN Certificate Guidelines and Limitations 1485
Managing Firewall Threat Defense Certificates 1486
Automatically Update CA Bundles 1487
Installing a Certificate Using Self-Signed Enrollment 1489
Installing a Certificate using EST Enrollment 1490
Installing a Certificate Using SCEP Enrollment 1491

Installing a Certificate Using Manual Enrollment 1491
Installing a Certificate Using a PKCS12 File 1492
Troubleshooting Firewall Threat Defense Certificates 1493
History for Certificates 1494

PART VI VPN 1495

CHAPTER 36 VPN Overview 1497
VPN Types 1497
VPN Basics 1498
Internet Key Exchange (IKE) 1499
IPsec 1499
VPN Packet Flow 1500
IPsec Flow Offload 1500
VPN Licensing 1501
How Secure Should a VPN Connection Be? 1502
Complying with Security Certification Requirements 1502
Deciding Which Encryption Algorithm to Use 1502
Deciding Which Hash Algorithms to Use 1503
Deciding Which Diffie-Hellman Modulus Group to Use 1503
Deciding Which Authentication Method to Use 1504
Pre-shared Keys 1504
PKI Infrastructure and Digital Certificates 1505
Removed or Deprecated Hash Algorithms, Encryption Algorithms, and Diffie-Hellman Modulus Groups 1506
VPN Topology Options 1507
Point-to-Point VPN Topology 1507
Hub and Spoke VPN Topology 1507
Full Mesh VPN Topology 1508
Implicit Topologies 1509

CHAPTER 37 Site-to-Site VPNs 1511
About Site-to-Site VPN 1511
Secure Firewall Threat Defense Site-to-site VPN Guidelines and Limitations 1513

Types of Site-to-Site VPN Topologies 1514
Requirements and Prerequisites for Site-to-Site VPN 1514
Manage Site-to-Site VPNs 1515
Configure a Policy-based Site-to-Site VPN 1516
Firewall Threat Defense VPN Endpoint Options 1517
Firewall Threat Defense VPN IKE Options 1521
Firewall Threat Defense VPN IPsec Options 1523
Firewall Threat Defense Advanced Site-to-site VPN Deployment Options 1526
Firewall Threat Defense VPN Advanced IKE Options 1526
Firewall Threat Defense VPN Advanced IPsec Options 1527
Firewall Threat Defense Advanced Site-to-site VPN Tunnel Options 1528
About Virtual Tunnel Interfaces 1529
Static VTI 1530
Dynamic VTI 1531
Guidelines and Limitations for Virtual Tunnel Interfaces 1533
Add a VTI Interface 1537
Create a Route-based Site-to-Site VPN 1538
Configure Endpoints for a Point to Point Topology 1539
Advanced Configurations for a Point to Point Topology in a Route-based VPN 1541
Configure Endpoints for a Hub and Spoke Topology 1542
Advanced Configurations for Hub and Spokes in a Route-based VPN 1544
Configure Multiple Hubs in a Route-based VPN 1545
Configure Routing for Multiple Hubs in a Route-based VPN 1547
Verify the Multiple Hubs Configuration in a Route-based VPN 1549
Route Traffic Through a Backup VTI Tunnel 1549
Configure Dynamic VTI for a Route-based Site-to-Site VPN 1551
How to Configure a Virtual Router with Dynamic VTI 1552
Configure Routing and AC Policies for VTI 1552
View Virtual Tunnel Information 1555
Deploy a SASE Tunnel on Umbrella 1555
Guidelines and Limitations for Configuring SASE Tunnels on Umbrella 1556
How to Deploy a SASE Tunnel on Umbrella 1557
Prerequisites for Configuring Umbrella SASE Tunnels 1558
Map Management Center Umbrella Parameters and Cisco Umbrella API Keys 1558

Configure a SASE Tunnel for Umbrella 1560
View SASE Tunnel Status 1561
Monitoring the Site-to-Site VPNs 1562
History for Site-to-Site VPN 1569

CHAPTER 38 **Remote Access VPN** 1573
Remote Access VPN Overview 1573
Remote Access VPN Features 1574
Secure Client Components 1576
Remote Access VPN Authentication 1576
Understanding Policy Enforcement of Permissions and Attributes 1578
Understanding AAA Server Connectivity 1578
License Requirements for Remote Access VPN 1579
Requirements and Prerequisites for Remote Access VPN 1579
Remote Access VPN Guidelines and Limitations 1580
Configuring a New Remote Access VPN Connection 1583
Prerequisites for Configuring Remote Access VPN 1584
Create a New Remote Access VPN Policy 1584
Update the Access Control Policy on the Secure Firewall Threat Defense Device 1587
(Optional) Configure NAT Exemption 1588
Configure DNS 1589
Add Secure Client Profile XML File 1589
(Optional) Configure Split Tunneling 1590
(Optional) Configure Dynamic Split Tunneling 1591
Verify Dynamic Split Tunneling Configuration 1592
Verify the Configuration 1592
Create a Copy of an Existing Remote Access VPN Policy 1593
Set Target Devices for a Remote Access VPN Policy 1593
Associate Local Realm with Remote Access VPN Policy 1594
Additional Remote Access VPN Configurations 1594
Configure Connection Profile Settings 1594
Configure IP Addresses for VPN Clients 1595
Configure AAA Settings for Remote Access VPN 1597
Create or Update Aliases for a Connection Profile 1612

Configure Access Interfaces for Remote Access VPN 1613
Configure Advanced Options for Remote Access VPN 1615
Cisco Secure Client Image 1615
Remote Access VPN Address Assignment Policy 1617
Configure Certificate to Connection Profile Mapping 1618
Configure Group Policies 1619
Configuring LDAP Attribute Mapping 1619
Configuring VPN Load Balancing 1621
Configure IPsec Settings 1624
Customize Cisco Secure Client 1629
Configure Secure Client Management VPN Tunnel 1640
Requirements and Prerequisites for Secure Client Management VPN Tunnel 1641
Limitations of Secure Client Management VPN Tunnel 1641
Configuring Secure Client Management VPN Tunnel on Firewall Threat Defense 1641
Multiple Certificate Authentication 1643
Guidelines and Limitations of Multiple Certificate Authentication 1644
Configuring Multiple Certificate Authentication 1644
Manage VPN Access of Remote Users Based on Geolocation 1645
Workflow to Manage VPN Access of Remote Users Based on Geolocation 1645
Guidelines and Limitations for Managing Remote Access VPN Users Based on Geolocation 1646
Monitor and Troubleshoot Service Access Policies 1646
Customizing Remote Access VPN AAA Settings 1650
Authenticate VPN Users via Client Certificates 1650
Configure VPN User Authentication via Client Certificate and AAA Server 1651
Manage Password Changes over VPN Sessions 1653
Send Accounting Records to the RADIUS Server 1654
Delegating Group Policy Selection to Authorization Server 1654
Override the Selection of Group Policy or Other Attributes by the Authorization Server 1655
Deny VPN Access to a User Group 1656
Restrict Connection Profile Selection for a User Group 1657
Update the Secure Client Profile for Remote Access VPN Clients 1658
RADIUS Dynamic Authorization 1658
Configuring RADIUS Dynamic Authorization 1659
Two-Factor Authentication 1660

Configuring RSA Two-Factor Authentication 1660
Configuring Duo Two-Factor Authentication 1661
Secondary Authentication 1663
Configure Remote Access VPN Secondary Authentication 1664
Single Sign-On Authentication with SAML 2.0 1666
Guidelines and Limitations for SAML 2.0 1666
Configuring a SAML Single Sign-On Authentication 1668
Configuring SAML Authorization 1669
Advanced Secure Client Configurations 1670
Configure Secure Client Modules on a Firewall Threat Defense 1670
Types of Secure Client Modules 1671
Prerequisites for Configuring Secure Client Modules 1672
Guidelines for Configuring Secure Client Modules 1673
Install Secure Client Modules using a Firewall Threat Defense 1674
Configure a Remote Access VPN Group Policy with Secure Client Modules 1674
Verify Secure Client Modules Configuration 1675
Configure Application-Based (Per App VPN) Remote Access VPN on Mobile Devices 1676
Prerequisites and Licensing for Configuring Per App VPN Tunnels 1676
Determine the Application IDs for Mobile Applications 1676
Configure Application-Based VPN Tunnels 1677
Verify Per App Configuration 1679
Remote Access VPN Examples 1679
How to Limit Secure Client Bandwidth Per User 1679
How to Use VPN Identity for User-Id Based Access Control Rules 1680
Configure Firewall Threat Defense Multiple Certificate Authentication 1680
History for Remote Access VPNs 1684

CHAPTER 39 **Dynamic Access Policies** 1689
About Secure Firewall Threat Defense Dynamic Access Policy 1689
Hierarchy of Policy Enforcement of Permissions and Attributes in Firewall Threat Defense 1689
Prerequisites for Dynamic Access Policy 1691
Guidelines and Limitations for Dynamic Access Policies 1691
Configure a Dynamic Access Policy (DAP) 1691
Create a Dynamic Access Policy 1691

Create a Dynamic Access Policy Record 1692
Configure Posture Assessment Criteria 1692
Configure AAA Criteria Settings for DAP 1693
Configure Endpoint Attribute Selection Criteria in DAP 1694
Add an Anti-Malware Endpoint Attribute to a DAP 1695
Add a Device Endpoint Attribute to a DAP 1696
Add Secure Client Endpoint Attributes to a DAP 1696
Add NAC Endpoint Attributes to a DAP 1697
Add an Application Attribute to a DAP 1697
Add a Personal Firewall Endpoint Attribute to a DAP 1697
Add an Operating System Endpoint Attribute to a DAP 1698
Add a Process Endpoint Attribute to a DAP 1698
Add a Registry Endpoint Attribute to a DAP 1698
Add a File Endpoint Attribute to a DAP 1699
Add Certificate Authentication Attributes to a DAP 1699
Configure Advanced Settings for DAP 1700
Associate Dynamic Access Policy with Remote Access VPN 1701
History for Dynamic Access Policy 1702

CHAPTER 40 **VPN Monitoring and Troubleshooting 1703**
VPN Summary Dashboard 1703
Remote Access VPN Dashboard 1703
SD-WAN Summary Dashboard 1705
Prerequisites for Using SD-WAN Summary Dashboard 1705
Monitor WAN Devices and Interfaces Using the SD-WAN Summary Dashboard 1708
Monitor Application Performance Metrics of WAN Interfaces Using the SD-WAN Summary Dashboard 1710
VPN Session and User Information 1710
Viewing Remote Access VPN Active Sessions 1710
Viewing Remote Access VPN User Activity 1710
Site to Site VPN Connection Event Monitoring 1710
View Site to Site VPN Connection Events 1711
VPN Troubleshooting 1712
System Messages 1712

VPN System Logs 1712
Debug Commands 1712
debug aaa 1714
debug crypto 1714
debug ldap 1717
debug ssl 1718
debug webvpn 1718

PART VII **Access Control** 1721

CHAPTER 41 **Access Control Overview** 1723
Introduction to Access Control 1723
Introduction to Rules 1724
Filtering Rules by Device 1724
Rule and Other Policy Warnings 1725
Access Control Policy Default Action 1726
Deep Inspection Using File and Intrusion Policies 1728
Access Control Traffic Handling with Intrusion and File Policies 1729
File and Intrusion Inspection Order 1730
Access Control Policy Inheritance 1731
Best Practices for Application Control 1732
Recommendations for Application Control 1732
Best Practices for Configuring Application Control 1735
Application Characteristics 1736
Application-Specific Notes and Limitations 1737
Best Practices for Access Control Rules 1738
General Best Practices for Access Control 1738
Best Practices for Ordering Rules 1739
Rule Preemption 1739
Rule Actions and Rule Order 1740
Application Rule Order 1741
URL Rule Order 1741
Best Practices for Simplifying and Focusing Rules 1742
Maximum Number of Access Control Rules and Intrusion Policies 1742

CHAPTER 42 **Access Control Policies** 1745

Access Control Policy Components 1745
System-Created Access Control Policies 1746
Requirements and Prerequisites for Access Control Policies 1746
Managing Access Control Policies 1747
Creating a Basic Access Control Policy 1748
Editing an Access Control Policy 1749
Locking an Access Control Policy 1751
Managing Access Control Policy Inheritance 1752
Choosing a Base Access Control Policy 1753
Inheriting Access Control Policy Settings from the Base Policy 1753
Locking Settings in Descendant Access Control Policies 1754
Requiring an Access Control Policy in a Domain 1754
Assigning Devices to an Access Control Policy 1755
Logging Settings for Access Control Policies 1755
Access Control Policy Advanced Settings 1756
Associating Other Policies with Access Control 1762
Identifying and Fixing Anomalies with Policy Analyzer & Optimizer 1763
Viewing Rule Hit Counts 1764
Analyzing Rule Conflicts and Warnings 1766
Searching for Rules 1768
History for Access Control Policies 1769

CHAPTER 43 **Access Control Rules** 1773

Introduction to Access Control Rules 1773
Access Control Rule Management 1775
Access Control Rule Components 1776
Access Control Rule Order 1777
Access Control Rule Actions 1778
Access Control Rule Monitor Action 1778
Access Control Rule Trust Action 1778
Access Control Rule Blocking Actions 1779
Access Control Rule Interactive Blocking Actions 1780

Access Control Rule Allow Action 1780
Requirements and Prerequisites for Access Control Rules 1781
Guidelines and Limitations for Access Control Rules 1782
Managing Access Control Rules 1783
Adding an Access Control Rule Category 1783
Create and Edit Access Control Rules 1783
Access Control Rule Conditions 1785
Enabling and Disabling Access Control Rules 1792
Copying Access Control Rules from One Access Control Policy to Another 1792
Moving Access Control Rules to a Prefilter Policy 1793
Positioning an Access Control Rule 1795
Adding Comments to an Access Control Rule 1796
Examples for Access Control Rules 1796
How to Control Access Using Security Zones 1796
How to Control Application Usage 1797
How to Block Threats 1798
History for Access Control Rules 1801

CHAPTER 44 **Cisco Secure Dynamic Attributes Connector 1803**
About the Cisco Secure Dynamic Attributes Connector 1803
How It Works 1804
History for the Cisco Secure Dynamic Attributes Connector 1805
System Requirements for the Cisco Secure Dynamic Attributes Connector 1806
Enable the Cisco Secure Dynamic Attributes Connector 1806
Configure Networks and Subnets for Docker Containers 1807
About the Dashboard 1809
Dashboard of an Unconfigured System 1810
Dashboard of a Configured System 1811
Add, Edit, or Delete Connectors 1813
Add, Edit, or Delete Dynamic Attributes Filters 1814
Create a Connector 1816
Amazon Web Services Connector—About User Permissions and Imported Data 1816
Create an AWS User with Minimal Permissions for the Cisco Secure Dynamic Attributes Connector 1817

Create an AWS Connector 1818
Amazon Web Services Security Groups Connector—About User Permissions and Imported Data 1819
Create an AWS User with Minimal Permissions for the Cisco Secure Dynamic Attributes Connector 1819
Create an AWS Security Groups Connector 1821
Create an AWS Service Tags Connector 1822
Azure Connector—About User Permissions and Imported Data 1823
Create an Azure User with Minimal Permissions for the Cisco Secure Dynamic Attributes Connector 1823
Create an Azure Connector 1825
Create an Azure Service Tags Connector 1826
Create a Cisco Cyber Vision Connector 1827
Create a Generic Text Connector 1828
Create a GitHub Connector 1830
Google Cloud Connector—About User Permissions and Imported Data 1830
Create a Google Cloud User with Minimal Permissions for the Cisco Secure Dynamic Attributes Connector 1831
Create a Google Cloud Connector 1832
Create an Office 365 Connector 1833
vCenter Connector—About User Permissions and Imported Data 1834
Create a vCenter User with Minimal Permissions for the Cisco Secure Dynamic Attributes Connector 1834
Create a vCenter Connector 1835
Create a Webex Connector 1838
Create a Zoom Connector 1839
Create Dynamic Attributes Filters 1839
Dynamic Attribute Filter Examples 1841
Manually Get a Certificate Authority (CA) Chain 1842
Use Dynamic Objects in Access Control Policies 1845
About Dynamic Objects in Access Control Rules 1845
Dynamic Attributes Rule Conditions 1846
Create Access Control Rules Using Dynamic Attributes Filters 1846
Disable the Cisco Secure Dynamic Attributes Connector 1847
Troubleshoot Using the Command Line 1848
Troubleshoot Using the Firewall Management Center 1850

Manually Get a Certificate Authority (CA) Chain 1850
Security Requirements 1853
Internet Access Requirements 1854
History for the Cisco Secure Dynamic Attributes Connector 1855

CHAPTER 45 URL Filtering 1857
URL Filtering Overview 1857
About URL Filtering with Category and Reputation 1858
URL Category and Reputation Descriptions 1859
URL Filtering Data from the Cisco Cloud 1859
Best Practices for URL Filtering 1859
Filtering HTTPS Traffic 1862
Use Categories in URL Filtering 1863
License Requirements for URL Filtering 1864
Requirements and Prerequisites for URL Filtering 1864
How to Configure URL Filtering with Category and Reputation 1865
Enable URL Filtering Using Category and Reputation 1866
URL Filtering Options 1866
Configuring URL Conditions 1868
Rules with URL Conditions 1869
URL Rule Order 1870
DNS Filtering: Identify URL Reputation and Category During DNS Lookup 1870
Enable DNS Filtering to Identify URLs During Domain Lookup 1870
DNS Filtering Limitations 1871
DNS Filtering and Events 1871
Manual URL Filtering 1871
Manual URL Filtering Options 1872
Supplement or Selectively Override Category and Reputation-Based URL Filtering 1872
Configure HTTP Response Pages 1873
Limitations to HTTP Response Pages 1873
Requirements and Prerequisites for HTTP Response Pages 1874
Choosing HTTP Response Pages 1875
Configure Interactive Blocking with HTTP Response Pages 1875
Configuring Interactive Blocking 1876

Setting the User Bypass Timeout for a Blocked Website 1876
Configure URL Filtering Health Monitors 1877
Dispute URL Category and Reputation 1877
If the URL Category Set Changes, Take Action 1878
URL Category and Reputation Changes: Effect on Events 1879
Troubleshoot URL Filtering 1880
History for URL Filtering 1882

CHAPTER 46 **Security Intelligence** 1885
About Security Intelligence 1885
Best Practices for Security Intelligence 1886
License Requirements for Security Intelligence 1886
Requirements and Prerequisites for Security Intelligence 1887
Security Intelligence Sources 1887
Configure Security Intelligence 1888
Security Intelligence Options 1890
Security Intelligence Categories 1891
Block List Icons 1893
Configuration Example: Security Intelligence Blocking 1893
Security Intelligence Monitoring 1894
Override Security Intelligence Blocking 1895
Troubleshooting Security Intelligence 1895
Security Intelligence Categories Are Missing from the Available Options List 1896
History for Security Intelligence Block Listing 1896

CHAPTER 47 **DNS Policies** 1897
DNS Policy Overview 1897
Cisco Umbrella DNS Policies 1898
DNS Policy Components 1898
License Requirements for DNS Policies 1899
Requirements and Prerequisites for DNS Policies 1899
Managing DNS and Umbrella DNS Policies 1900
Creating Basic DNS Policies 1900
Editing DNS Policies 1901

DNS Rules 1901
Creating and Editing DNS Rules 1902
DNS Rule Management 1903
Enabling and Disabling DNS Rules 1903
DNS Rule Order Evaluation 1903
DNS Rule Actions 1904
DNS Rule Conditions 1905
Security Zone Rule Conditions 1905
Network Rule Conditions 1906
VLAN Tags Rule Conditions 1906
DNS Policy Rule Conditions 1906
How to Create DNS Rules 1907
Controlling Traffic Based on DNS and Security Zone 1907
Controlling Traffic Based on DNS and Network 1907
Controlling Traffic Based on DNS and VLAN 1908
Controlling Traffic Based on DNS List or Feed 1909
DNS Policy Deploy 1909
Cisco Umbrella DNS Policies 1909
How to Redirect DNS Requests to Cisco Umbrella 1910
Prerequisites for Configuring the Umbrella DNS Connector 1911
Configure Cisco Umbrella Connection Settings 1912
Create an Umbrella DNS Policy 1913
Edit Umbrella DNS Policies and Rules 1913
Associate the Umbrella DNS Policy with an Access Control Policy 1914

CHAPTER 48 **Prefiltering and Prefilter Policies** 1915
About Prefiltering 1915
About Prefilter Policies 1915
Tunnel vs Prefilter Rules 1916
Prefiltering vs Access Control 1917
Passthrough Tunnels and Access Control 1919
Best Practices for Fastpath Prefiltering 1920
Best Practices for Encapsulated Traffic Handling 1920
Requirements and Prerequisites for Prefilter Policies 1921

Configure Prefiltering 1922
Tunnel and Prefilter Rule Components 1923
Prefilter Rule Conditions 1925
Interface Rule Conditions 1925
Network Rule Conditions 1925
VLAN Tags Rule Conditions 1926
Port Rule Conditions for Prefilter Rules 1926
Time and Day Rule Conditions 1927
Tunnel Rule Conditions 1927
Encapsulation Rule Conditions 1927
Tunnel Zones and Prefiltering 1928
Using Tunnel Zones 1928
Creating Tunnel Zones 1930
Moving Prefilter Rules to an Access Control Policy 1931
Prefilter Policy Hit Counts 1933
Large Flow Offloads 1933
Flow Offload Limitations 1934
History for Prefiltering 1936

CHAPTER 49 **Service Policies** 1937
About Threat Defense Service Policies 1937
How Service Policies Relate to FlexConfig and Other Features 1938
What Are Connection Settings? 1938
Requirements and Prerequisites for Service Policies 1939
Guidelines and Limitations for Service Policies 1939
Configure Threat Defense Service Policies 1940
Configure a Service Policy Rule 1941
Bypass TCP State Checks for Asymetrical Routing (TCP State Bypass) 1943
The Asymetrical Routing Problem 1944
Guidelines and Limitations for TCP State Bypass 1945
Configure TCP State Bypass 1945
Disable TCP Sequence Randomization 1947
Examples for Service Policy Rules 1948
Protect Servers from a SYN Flood DoS Attack (TCP Intercept) 1948

Make the Firewall Threat Defense Device Appear on Traceroutes 1951
Monitoring Service Policies 1953
History for Threat Defense Service Policy 1954

CHAPTER 50 Threat Detection 1955
Portscan Detection and Prevention 1955
Pre-Defined Sensitivity Levels for Portscan Detection 1955
Detection in the Low Sensitivity Level 1957
Best Practices for Portscan Prevention 1957
Requirements and Prerequisites for Threat Detection 1958
Guidelines and Limitations for Threat Detection 1958
Configure Portscan Detection and Prevention 1959
Monitoring Threat Detection 1961
Viewing Portscan Alerts 1961
Monitoring Portscan on the Firewall 1962
Unblocking A Host 1963
History for Threat Detection 1963

CHAPTER 51 Intelligent Application Bypass 1965
Introduction to IAB 1965
IAB Options 1966
Requirements and Prerequisites for Intelligent Application Bypass 1968
Configuring Intelligent Application Bypass 1968
IAB Logging and Analysis 1969

CHAPTER 52 Content Restriction 1973
About Content Restriction 1973
Requirements and Prerequisites for Content Restriction 1974
Guidelines and Limitations for Content Restriction 1975
Using Access Control Rules to Enforce Content Restriction 1975
Safe Search Options for Access Control Rules 1976
Using a DNS Sinkhole to Enforce Content Restriction 1976

CHAPTER 53 Zero Trust Access 1979

About Clientless and Universal Zero Trust Network Access 1979
Clientless Zero Trust Access 1979
How Threat Defense Works with Zero Trust Access 1980
Why Use Zero Trust Access? 1981
Components of a Zero Trust Access Configuration 1981
Zero Trust Access Workflow 1982
Limitations for Zero Trust Network Access 1983
Prerequisites for Zero Trust Application Policy 1984
Manage Zero Trust Application Policies 1985
Create a Zero Trust Application Policy 1985
Create an Application Group 1986
Create an Application 1988
Set Targeted Devices for Zero Trust Access Policy 1989
Edit a Zero Trust Application Policy 1990
Monitor Zero Trust Sessions 1992
Overview of Universal Zero Trust Network Access 1994
How Threat Defense Works with Universal ZTNA 1995
Prerequisites for Universal Zero Trust Network Access 1995
Limitations of Universal Zero Trust Network Access 1996
Integrate Firewall Management Center with Security Cloud Control 1997
Security Cloud Control Settings 2000
Configure Security Devices 2001
History for Zero Trust Access 2007

PART VIII **SD-WAN** 2009

CHAPTER 54 **SD-WAN Capabilities** 2011
Overview of SD-WAN Capabilities 2011
Features 2012
Using SD-WAN Wizard for Secure Branch Network Deployment 2013
Guidelines and Limitations for Using SD-WAN Wizard 2014
Prerequisites for Using the SD-WAN Wizard 2015
Configure an SD-WAN Topology Using the SD-WAN Wizard 2016
Add a Dynamic Virtual Tunnel Interface for a Hub 2019

Sample Configurations for Dual ISP Deployment Using SD-WAN Wizard 2020
Dual ISP Deployment: Two Hubs and Four Spokes in the Same Region 2020
Dual ISP Deployment: Two Hubs and Four Spokes in Different Regions 2022
Verify Tunnel Statuses of an SD-WAN Topology 2025
Use Cases for SD-WAN Capabilities 2028

PART IX **Intrusion Detection and Prevention** 2029

CHAPTER 55 **An Overview of Network Analysis and Intrusion Policies** 2031
About Network Analysis and Intrusion Policies 2031
Snort Inspection Engine 2032
Snort 3 2033
Guidelines and Limitations for Network Analysis and Intrusion Policies 2035
How Policies Examine Traffic For Intrusions 2036
Decoding, Normalizing, and Preprocessing: Network Analysis Policies 2037
Access Control Rules: Intrusion Policy Selection 2038
Intrusion Inspection: Intrusion Policies, Rules, and Variable Sets 2039
Intrusion Event Generation 2041
Asymmetrical Flow Inspection in Snort 2041
System-Provided and Custom Network Analysis and Intrusion Policies 2042
System-Provided Network Analysis and Intrusion Policies 2043
Benefits of Custom Network Analysis and Intrusion Policies 2044
Benefits of Custom Network Analysis Policies 2045
Benefits of Custom Intrusion Policies 2045
Limitations of Custom Policies 2046
Prerequisites for Network Analysis and Intrusion Policies 2048

CHAPTER 56 **Migrate from Snort 2 to Snort 3** 2049
Snort 3 Inspection Engine 2049
Prerequisites for Network Analysis and Intrusion Policies 2050
How to Migrate from Snort 2 to Snort 3 2050
Prerequisites for Migrating from Snort 2 to Snort 3 2050
Enable Snort 3 on an Individual Device 2051
Enable Snort 3 on Multiple Devices 2051

Convert Snort 2 Custom IPS Rules to Snort 3 2052
Convert all Snort 2 Custom Rules across all Intrusion Policies to Snort 3 2052
Convert Snort 2 Custom Rules of a Single Intrusion Policy to Snort 3 2053
View Snort 2 and Snort 3 Base Policy Mapping 2054
Synchronize Snort 2 Rules with Snort 3 2054
Deploy Configuration Changes 2055

CHAPTER 57 **Get Started with Snort 3 Intrusion Policies** 2059
Overview of Intrusion Policies 2059
Prerequisites for Network Analysis and Intrusion Policies 2060
Create a Custom Snort 3 Intrusion Policy 2060
Edit Snort 3 Intrusion Policies 2061
Rule Group Reporting 2065
Rule Action Logging 2065
Change the Base Policy of an Intrusion Policy 2066
View Snort 2 and Snort 3 Base Policy Mapping 2066
Synchronize Snort 2 Rules with Snort 3 2067
Manage Intrusion Policies 2068
Access Control Rule Configuration to Perform Intrusion Prevention 2069
Access Control Rule Configuration and Intrusion Policies 2070
Configure an Access Control Rule to Perform Intrusion Prevention 2070
Deploy Configuration Changes 2071

CHAPTER 58 **Tune Intrusion Policies Using Rules** 2073
Overview of Tuning Intrusion Rules 2073
Intrusion Rule Types 2074
Prerequisites for Network Analysis and Intrusion Policies 2075
Custom Rules in Snort 3 2075
View Snort 3 Intrusion Rules in an Intrusion Policy 2078
Intrusion Rule Action 2079
Intrusion Rule Action Options 2079
Set Intrusion Rule Action 2079
Intrusion Event Notification Filters in an Intrusion Policy 2080
Intrusion Event Thresholds 2080

Set Intrusion Event Thresholds 2080
Set Threshold for an Intrusion Rule in Snort 3 2082
View and Delete Intrusion Event Thresholds 2083
Intrusion Policy Suppression Configuration 2083
Intrusion Policy Suppression Types 2083
Set Suppression for an Intrusion Rule in Snort 3 2084
View and Delete Suppression Conditions 2084
Add Intrusion Rule Comments 2085
Snort 2 Custom Rules Conversion to Snort 3 2085
Convert all Snort 2 Custom Rules across all Intrusion Policies to Snort 3 2086
Convert Snort 2 Custom Rules of a Single Intrusion Policy to Snort 3 2087
Add Custom Rules to Rule Groups 2087
Add Rule Groups with Custom Rules to an Intrusion Policy 2088
Manage Custom Rules in Snort 3 2089
Delete Custom Rules 2090
Delete Rule Groups 2090

CHAPTER 59 **Tailor Intrusion Protection for Your Network Assets** 2093
Snort 3 Rule Changes in LSP Updates 2093
Overview of Secure Firewall Recommended Rules 2093
Prerequisites for Network Analysis and Intrusion Policies 2094
Generate New Secure Firewall Recommendations in Snort 3 2095

CHAPTER 60 **Intrusion Prevention Performance Tuning** 2099
About Intrusion Prevention Performance Tuning 2099
License Requirements for Intrusion Prevention Performance Tuning 2100
Requirements and Prerequisites for Intrusion Prevention Performance Tuning 2100
Limiting Pattern Matching for Intrusions 2100
Regular Expression Limits Overrides for Intrusion Rules 2101
Overriding Regular Expression Limits for Intrusion Rules 2102
Per Packet Intrusion Event Generation Limits 2103
Limiting Intrusion Events Generated Per Packet 2103
Packet and Intrusion Rule Latency Threshold Configuration 2104
Latency-Based Performance Settings 2104

Packet Latency Thresholding 2104
Packet Latency Thresholding Notes 2105
Enabling Packet Latency Thresholding 2106
Configuring Packet Latency Thresholding 2106
Rule Latency Thresholding 2107
Rule Latency Thresholding Notes 2108
Configuring Rule Latency Thresholding 2109
Intrusion Performance Statistic Logging Configuration 2110
Configuring Intrusion Performance Statistic Logging 2110

CHAPTER 61 **Get Started with Snort 3 Network Analysis Policies** 2113
Overview of Network Analysis Policies 2113
Manage Network Analysis Policies 2114
Snort 3 Definitions and Terminologies for Network Analysis Policy 2115
Prerequisites for Network Analysis and Intrusion Policies 2117
Custom Network Analysis Policy Creation for Snort 3 2117
Common Industrial Protocol Safety 2121
Detect and Block Safety Segments in CIP Packets 2122
Network Analysis Policy Mapping 2122
View Network Analysis Policy Mapping 2123
Create a Network Analysis Policy 2123
Modify the Network Analysis Policy 2123
Search for an Inspector on the Network Analysis Policy Page 2124
Copy the Inspector Configuration 2125
Customize the Network Analysis Policy 2125
Make Inline Edit for an Inspector to Override Configuration 2129
Revert Unsaved Changes during Inline Edits 2130
View the List of Inspectors with Overrides 2130
Revert Overridden Configuration to Default Configuration 2131
Validate Snort 3 Policies 2131
Examples of Custom Network Analysis Policy Configuration 2134
Network Analysis Policy Settings and Cached Changes 2145

CHAPTER 62 **Encrypted Visibility Engine** 2147

Overview of Encrypted Visibility Engine 2147
How EVE Works 2148
Indications of Compromise Events 2149
QUIC Fingerprinting in EVE 2149
Configure EVE 2149
View Encrypted Visibility Engine Events 2151
View EVE Dashboard 2151
Configure EVE Exception Rules 2152
Add Exception Rule from Unified Events 2153
Event Enrichment 2154

CHAPTER 63 **Elephant Flow Detection** 2155
About Elephant Flow Detection and Remediation 2155
Elephant Flow Upgrade from Intelligent Application Bypass 2155
Configure Elephant Flow 2156

CHAPTER 64 **Use Case - Migrate from Snort 2 to Snort 3 In Secure Firewall Management Center** 2159
Migrate from Snort 2 to Snort 3 2159
Benefits of Migrating to Snort 3 2159
Sample Business Scenario 2160
Best Practices for Migrating from Snort 2 to Snort 3 2160
Prerequisites 2160
End-to-End Migration Workflow 2160
Enable Snort 3 on Threat Defense 2161
Convert Snort 2 Rules of a Single Intrusion Policy to Snort 3 2162
Deploy Configuration Changes 2167

CHAPTER 65 **Use Case - Generate Snort 3 Recommendations In Secure Firewall Management Center** 2171
Snort 3 Rule Recommendations 2171
Benefits 2172
Sample Business Scenario 2172
Best Practices 2172
Prerequisites 2172
Generate Snort 3 Recommendations 2172

Deploy Configuration Changes 2175

CHAPTER 66 Use Case - Block Traffic Based on the EVE Threat Confidence Score 2179

About Encrypted Visibility Engine 2179

Benefits 2179

Sample Business Scenario 2179

Prerequisites 2180

High-Level Workflow 2180

Configure Block Thresholds in EVE 2180

View EVE Events 2184

Additional References 2185

CHAPTER 67 Use Case - Configure Elephant Flow Detection Outcomes 2187

About Elephant Flows 2187

Benefits of Elephant Flow Detection and Remediation 2187

Elephant Flow Workflow 2188

Sample Business Scenario 2188

Prerequisites 2189

Configure Elephant Flow Parameters 2189

View Events for Elephant Flows 2192

Configure Elephant Flow Remediation Exemption 2193

View Events for Elephant Flow Remediation Exemption 2196

Additional References 2196

CHAPTER 68 Mitigate Threats Using MITRE Framework in Snort 3 Intrusion Policies 2197

About MITRE ATT&CK Framework 2197

Benefits of MITRE Framework 2198

Sample Business Scenario for MITRE Network 2198

Prerequisites for MITRE Framework 2198

View and Edit Your Snort 3 Intrusion Policy 2198

View Intrusion Events 2204

Additional References 2206

PART X Network Malware Protection and File Policies 2207

CHAPTER 69 **Network Malware Protection and File Policies** 2209

About Network Malware Protection and File Policies 2209

File Policies 2210

Requirements and Prerequisites for File Policies 2210

License Requirements for File and Malware Policies 2211

Best Practices for File Policies and Malware Detection 2211

File Rule Best Practices 2211

File Detection Best Practices 2212

File Blocking Best Practices 2212

File Policy Best Practices 2213

How to Configure Malware Protection 2214

Plan and Prepare for Malware Protection 2215

Configure File Policies 2216

Add File Policies to Your Access Control Configuration 2216

Configuring an Access Control Rule to Perform Malware Protection 2217

Set Up Maintenance and Monitoring of Malware Protection 2218

Cloud Connections for Malware Protection 2219

AMP Cloud Connection Configurations 2220

Requirements and Best Practices for AMP Cloud Connections 2220

Choose an AMP Cloud 2221

Cisco AMP Private Cloud 2221

Managing Connections to the AMP Cloud (Public or Private) 2223

Change AMP Options 2224

Dynamic Analysis Connections 2224

Requirements for Dynamic Analysis 2224

Viewing the Default Dynamic Analysis Connection 2224

Dynamic Analysis On-Premises Appliance (Cisco Secure Malware Analytics) 2225

Enabling Access to Dynamic Analysis Results in the Public Cloud 2226

Maintain Your System: File Types Eligible for Dynamic Analysis 2227

File Policies and File Rules 2228

Create or Edit a File Policy 2228

Advanced and Archive File Inspection Options 2229

Managing File Policies 2232

File Rules 2233
File Rule Components 2233
File Rule Actions 2234
Creating File Rules 2241
Access Control Rule Logging for Malware Protection 2243
Retrospective Disposition Changes 2243
File and Malware Inspection Performance and Storage Options 2243
Tuning File and Malware Inspection Performance and Storage 2244
(Optional) Malware Protection with Secure Endpoint 2245
Comparison of Malware Protection: Firepower vs. Secure Endpoint 2246
About Integrating Firepower with Secure Endpoint 2246
Benefits of Integrating Firepower and Secure Endpoint 2247
Secure Endpoint and AMP Private Cloud 2247
Integrate Firepower and Secure Endpoint 2247
History for Network Malware Protection and File Policies 2250

PART XI **Encrypted Traffic Handling** 2251

CHAPTER 70 **Traffic Decryption Overview** 2253
Traffic Decryption Explained 2253
TLS/SSL Handshake Processing 2255
ClientHello Message Handling 2255
ServerHello and Server Certificate Message Handling 2258
Decryption Rule and Policy Basics 2260
The Case for Decryption 2261
When to Decrypt Traffic, When Not to Decrypt 2261
Decrypt and Resign (Outgoing Traffic) 2263
Known Key Decryption (Incoming Traffic) 2263
Other Decryption Rule Actions 2264
Decryption Rule Components 2264
Decryption Rule Order Evaluation 2265
Multi-Rule Example 2266
TLS Crypto Acceleration 2268
TLS Crypto Acceleration Guidelines and Limitations 2269

View the Status of TLS Crypto Acceleration 2270
DTLS Crypto Acceleration 2271
Prerequisites for DTLS Crypto Acceleration 2271
Monitoring DTLS Crypto Acceleration 2271
How to Configure Decryption Policies and Rules 2272
History for Decryption Policy 2273

CHAPTER 71 Decryption Policies 2279
About Decryption Policies 2279
Requirements and Prerequisites for Decryption Policies 2281
Create a Decryption Policy 2281
Create a Create a Decryption Policy with Outbound Connection Protection 2281
Create a Create a Decryption Policy with Inbound Connection Protection 2285
Decryption Policy Block Connections 2287
Decryption Policy Exclusions 2287
Generate an Internal CA for Outbound Protection 2291
Upload an Internal CA for Outbound Protection 2292
Upload an Internal Certificate for Inbound Protection 2293
Create a Create a Decryption Policy with Other Rule Actions 2294
Decryption Policy Default Actions 2295
Default Handling Options for Undecryptable Traffic 2296
Set Default Handling for Undecryptable Traffic 2297
Decryption Policy Advanced Options 2298
TLS 1.3 Decryption Best Practices 2300

CHAPTER 72 Decryption Rules 2303
Decryption Rules Overview 2303
Requirements and Prerequisites for Decryption Rules 2303
Decryption Rule Guidelines and Limitations 2304
Guidelines for Using TLS/SSL Decryption 2304
Decryption Rule Unsupported Features 2305
TLS/SSL Do Not Decrypt Guidelines 2306
TLS/SSL Decrypt - Resign Guidelines 2307
TLS/SSL Decrypt - Known Key Guidelines 2309

TLS/SSL Block Guidelines 2310
TLS/SSL Certificate Pinning Guidelines 2310
TLS/SSL Heartbeat Guidelines 2311
TLS/SSL Anonymous Cipher Suite Limitation 2311
TLS/SSL Normalizer Guidelines 2311
Other Decryption Rule Guidelines 2311
Decryption Rule Traffic Handling 2312
Encrypted Traffic Inspection Configuration 2314
Decryption Rule Order Evaluation 2315
Decryption Rule Conditions 2316
Security Zone Rule Conditions 2317
Security Zone Conditions and Multitenancy 2317
Network Rule Conditions 2318
VLAN Tags Rule Conditions 2318
User Rule Conditions 2318
Client Threat Rule Conditions 2319
Application Rule Conditions 2319
Port Rule Conditions 2321
Category Rule Conditions 2321
Server Certificate-Based Decryption Rule Conditions 2321
Certificate Decryption Rule Conditions 2322
Distinguished Name (DN) Rule Conditions 2323
Trusting External Certificate Authorities 2327
Certificate Status Decryption Rule Conditions 2329
Cipher Suite Decryption Rule Conditions 2331
Encryption Protocol Version Decryption Rule Conditions 2334
Decryption Rule Actions 2334
Decryption Rule Monitor Action 2334
Decryption Rule Do Not Decrypt Action 2335
Decryption Rule Blocking Actions 2336
Decryption Rule Decrypt Actions 2337
Select Internal Certificate Objects 2337
Monitor TLS/SSL Hardware Acceleration 2337
Informational Counters 2337

Alert Counters 2338
Error Counters 2338
Fatal Counters 2339
Troubleshoot Decryption Rules 2339
About TLS/SSL Oversubscription 2339
Troubleshoot TLS/SSL Oversubscription 2340
About TLS Heartbeat 2341
Troubleshoot TLS Heartbeat 2342
About TLS/SSL Pinning 2343
Troubleshoot TLS/SSL Pinning 2344
Troubleshoot Unknown or Bad Certificates or Certificate Authorities 2346
Verify TLS/SSL Cipher Suites 2348
Troubleshooting Using Crypto Archives 2349

CHAPTER 73 **Decryption Rules and Policy Example** 2351
Decryption Rule Examples 2351
Run the Decryption Policy Wizard 2351
Decryption Policy Exclusions 2353
First Manual Do Not Decrypt Rule: Specific Traffic 2356
Next Manual Rule: Decrypt Specific Test Traffic 2358
Last Manual Decryption Rules: Block or Monitor Certificates and Protocol Versions 2359
Example: Decryption Rule to Monitor or Block Certificate Status 2361
Example: Decryption Rule to Monitor or Block Protocol Versions 2363
Optional Example: Manual Decryption Rule to Monitor or Block Certificate Distinguished Name 2364
Associate the Decryption Policy with an Access Control Policy and Advanced Settings 2366
Traffic to Prefilter 2368
Decryption Rule Settings 2368

PART XII **User Identity** 2369

CHAPTER 74 **User Identity Overview** 2371
About User Identity 2371
Identity Terminology 2372
About User Identity Sources 2372

Best Practices for User Identity 2374
Identity Deployments 2376
How to Set Up an Identity Policy 2378
The User Activity Database 2381
The Users Database 2382
Host and User Limits 2383
Host Limit 2383
User Limits for Microsoft Active Directory 2384
User Limits for Microsoft Azure Active Directory Realms 2386
Identity Realm Limit 2387
User Limits for Microsoft Azure Active Directory Realms 2388

CHAPTER 75 Realms 2391
License Requirements for Realms 2391
Requirements and Prerequisites for Realms 2391
Create a Microsoft Azure AD (SAML) Realm 2392
How to Create a Microsoft Azure AD (SAML) for Passive Authentication 2393
About Entra ID and Cisco ISE with Resource Owned Password Credentials 2394
About Entra ID and Cisco ISE with TEAP/EAP-TLS 2394
How to Configure ISE for Microsoft Azure AD (SAML) 2395
Configure Microsoft Entra ID for Passive Authentication 2395
Configure Entra ID Basic Settings 2397
Get Required Information For Your Microsoft Azure AD (SAML) Realm 2398
Create a Microsoft Azure AD (SAML) Realm for Passive Authentication 2400
How to Create a Microsoft Azure AD (SAML) Realm for Active Authentication (Captive Portal) 2403
Configure Entra ID Basic Settings 2406
Configure a Single Sign-On (SSO) App in Entra ID 2407
Create a Decryption Rule with Decrypt - Resign Action 2408
Get Required Information For Your Microsoft Azure AD (SAML) Realm (Active Authentication Only) 2409
Create a Microsoft Azure AD (SAML) Realm for Active Authentication (Captive Portal) 2412
Create a Microsoft Azure AD (SAML) Realm for Passive Authentication 2417
About Entra ID and Cisco ISE with Resource Owned Password Credentials 2418
About Entra ID and Cisco ISE with TEAP/EAP-TLS 2419

How to Create a Microsoft Azure AD (SAML) for Passive Authentication 2419
Configure Microsoft Entra ID for Passive Authentication 2420
How to Configure ISE for Microsoft Azure AD (SAML) 2421
Get Required Information For Your Microsoft Azure AD (SAML) Realm 2422
Create an Azure AD Realm 2424
Azure AD (SAML) User Session Timeout 2426
Create an LDAP Realm or an Active Directory Realm and Realm Directory 2426
About Realms and Realm Sequences 2429
Realms and Trusted Domains 2431
Supported Servers for Realms 2434
Supported Server Object Class and Attribute Names 2435
Prerequisites for Kerberos Authentication 2436
Realm Fields 2436
Realm Directory and Synchronize fields 2440
Connect Securely to Active Directory or LDAP 2442
Find the Active Directory Server's Name 2442
Export the Active Directory Server's Root Certificate 2443
Synchronize Users and Groups 2445
Create a Realm Sequence 2446
Configure the Firewall Management Center for Cross-Domain-Trust: The Setup 2447
Configure the Secure Firewall Management Center for Cross-Domain-Trust Step 1: Configure Realms and Directories 2448
Configure the Secure Firewall Management Center for Cross-Domain-Trust Step 2: Synchronize Users and Groups 2453
Configure the Secure Firewall Management Center for Cross-Domain-Trust Step 3: Resolve Issues 2454
Manage a Realm 2455
Compare Realms 2456
Troubleshoot Realms and User Downloads 2456
Detect Realm or User Mismatches 2460
Troubleshoot Cross-Domain Trust 2461
History for Realms 2464

CHAPTER 76 **User Control with the Passive Identity Agent** 2467
The Passive Identity Agent Identity Source 2467

Deploy the Passive Identity Agent 2469
Simple Passive Identity Agent Deployment 2469
Single Passive Identity Agent Monitoring Multiple Domain Controllers 2470
Multiple Passive Identity Agents Monitoring Multiple Domain Controllers 2471
Passive Identity Agent Primary/Secondary Agent Deployments 2473
How to Create a Passive Identity Agent Identity Source 2474
Configure the Passive Identity Agent 2476
Create a Microsoft Active Directory Realm 2476
Create a Passive Identity Agent Identity Source 2476
Create a Standalone Passive Identity Agent Identity Source 2477
Create a Primary or Secondary Passive Identity Agent Identity Source 2479
About Passive Identity Agent Roles 2481
Create a Secure Firewall Management Center User for the Passive Identity Agent 2482
Troubleshoot the Passive Identity Agent 2484
About Passive Identity Agent Installation 2485
Prerequisites to Installing the Passive Identity Agent 2485
Install the Passive Identity Agent Software 2487
Uninstall the Passive Identity Agent Software 2491
Upgrade the Passive Identity Agent Software 2491
Monitor the Passive Identity Agent 2491
Manage the Passive Identity Agent 2493
Edit Passive Identity Agents 2493
Delete a Standalone Passive Identity Agent 2493
Delete Primary and Secondary Passive Identity Agents 2494
Troubleshoot the Passive Identity Agent 2494
Security Requirements for the Passive Identity Agent 2495
Internet Access Requirements for the Passive Identity Agent 2496
History for the Passive Identity Agent 2497

CHAPTER 77 **User Control with ISE/ISE-PIC** 2499
The ISE/ISE-PIC Identity Source 2499
Source and Destination Security Group Tag (SGT) Matching 2500
License Requirements for ISE/ISE-PIC 2501
Requirements and Prerequisites for ISE/ISE-PIC 2501

ISE/ISE-PIC Guidelines and Limitations 2501
How to Configure ISE/ISE-PIC for User Control 2504
How to Configure ISE/ISE-PIC Without a Realm 2504
How to Configure ISE/ISE-PIC for User Control Using a Realm 2505
Configure ISE/ISE-PIC 2506
Configure Security Groups and SXP Publishing in ISE 2506
Export Certificates from the ISE/ISE-PIC Server for Use in the Firewall Management Center 2509
Export a System Certificate 2510
Generate a Self-Signed Certificate 2510
Import ISE/ISE-PIC Certificates 2511
Ways to Configure the Cisco Identity Services Engine (Cisco ISE) Identity Source 2512
About Cisco ISE Quick Configuration 2512
Prerequisites for ISE Quick Configuration 2513
Quick Configuration 2514
Cisco Identity Services Engine (Cisco ISE) Quick Configuration Results 2515
Cisco ISE Advanced Configuration 2518
ISE/ISE-PIC Configuration Fields 2520
Troubleshoot the ISE/ISE-PIC or Cisco TrustSec Issues 2521
History for ISE/ISE-PIC 2523

CHAPTER 78 **User Control with Captive Portal** 2525
The Captive Portal Identity Source 2525
About Hostname Redirect 2526
License Requirements for Captive Portal 2526
Requirements and Prerequisites for Captive Portal 2526
Captive Portal Guidelines and Limitations 2526
How to Configure the Captive Portal for User Control 2529
Configure the Captive Portal Part 1: Create a Network Object 2531
Configure the Captive Portal Part 2: Create an Identity Policy and Active Authentication Rule 2532
Update a Custom Authentication Form 2534
Configure the Captive Portal Part 3: Create a TCP Port Access Control Rule 2534
Configure the Captive Portal Part 4: Create a User Access Control Rule 2535
Captive Portal Example: Create a Decryption Policy with an Outbound Rule 2536

Configure Captive Portal Part 6: Associate Identity and Decryption Policies with the Access Control Policy 2538
Captive Portal Fields 2539
Exclude Applications from Captive Portal 2540
Troubleshoot the Captive Portal Identity Source 2541
History for Captive Portal 2543

CHAPTER 79 **User Control with Remote Access VPN 2545**
The Remote Access VPN Identity Source 2545
Configure RA VPN for User Control 2546
Troubleshoot the Remote Access VPN Identity Source 2546
Not Observing Correct Settings for VPN Statistics 2547
History for RA VPN 2548

CHAPTER 80 **User Control with TS Agent 2549**
The Terminal Services (TS) Agent Identity Source 2549
TS Agent Guidelines 2550
User Control with TS Agent 2550
Troubleshoot the TS Agent Identity Source 2550
History for TS Agent 2551

CHAPTER 81 **User Identity Policies 2553**
About Identity Policies 2553
License Requirements for Identity Policies 2554
Requirements and Prerequisites for Identity Policies 2554
Create an Identity Policy 2555
Create an Identity Mapping Filter 2556
Identity Rule Conditions 2557
Security Zone Rule Conditions 2557
Security Zone Conditions and Multitenancy 2558
Network Rule Conditions 2558
Redirect to Host Name Network Rule Conditions 2558
VLAN Tags Rule Conditions 2559
Port Rule Conditions 2559

Port, Protocol, and ICMP Code Rule Conditions 2560
Realm & Settings Rule Conditions 2561
Create an Identity Rule 2563
Identity Rule Fields 2564
Sample Identity Policies and Rules 2565
Create an Identity Policy with a Passive Authentication Rule 2566
Create a Sample Identity Policy with an Active Authentication Rule 2568
Active Authentication Using a Realm 2570
Active Authentication Using a Realm Sequence 2571
Manage an Identity Policy 2572
Manage an Identity Rule 2573
Troubleshoot User Control 2573

PART XIII **Network Discovery** 2577

CHAPTER 82 **Network Discovery Overview** 2579
About Detection of Host, Application, and User Data 2579
Host and Application Detection Fundamentals 2580
Passive Detection of Operating System and Host Data 2580
Active Detection of Operating System and Host Data 2580
Current Identities for Applications and Operating Systems 2581
Current User Identities 2582
Application and Operating System Identity Conflicts 2583
NetFlow Data 2583
Requirements for Using NetFlow Data 2584
Differences between NetFlow and Managed Device Data 2584

CHAPTER 83 **Host Identity Sources** 2587
Overview: Host Data Collection 2587
Requirements and Prerequisites for Host Identity Sources 2588
Determining Which Host Operating Systems the System Can Detect 2588
Identifying Host Operating Systems 2588
Custom Fingerprinting 2589
Managing Fingerprints 2590

Activating and Deactivating Fingerprints 2590
Editing an Active Fingerprint 2591
Editing an Inactive Fingerprint 2591
Creating a Custom Fingerprint for Clients 2592
Creating a Custom Fingerprint for Servers 2594
Host Input Data 2597
Requirements for Using Third-Party Data 2597
Third-Party Product Mappings 2598
Mapping Third-Party Products 2598
Mapping Third-Party Product Fixes 2599
Mapping Third-Party Vulnerabilities 2600
Custom Product Mappings 2601
Creating Custom Product Mappings 2601
Editing Custom Product Mapping Lists 2602
Activating and Deactivating Custom Product Mappings 2603
Configuring the Host Input Client 2603
Nmap Scanning 2604
Nmap Remediation Options 2605
Nmap Scanning Guidelines 2609
Example: Using Nmap to Resolve Unknown Operating Systems 2610
Example: Using Nmap to Respond to New Hosts 2611
Managing Nmap Scanning 2612
Adding an Nmap Scan Instance 2613
Editing an Nmap Scan Instance 2614
Adding an Nmap Scan Target 2615
Editing an Nmap Scan Target 2616
Creating an Nmap Remediation 2616
Editing an Nmap Remediation 2618
Running an On-Demand Nmap Scan 2619
Nmap Scan Results 2619
Viewing Nmap Scan Results 2620
Nmap Scan Results Fields 2621
Importing Nmap Scan Results 2621
History for Host Identity Sources 2622

CHAPTER 84 **Application Detection** 2623

Overview: Application Detection 2623

Application Detector Fundamentals 2624

Identification of Application Protocols in the Web Interface 2625

Implied Application Protocol Detection from Client Detection 2626

Host Limits and Discovery Event Logging 2626

Special Considerations for Application Detection 2627

Application Detection in Snort 3 2628

Requirements and Prerequisites for Application Detection 2629

Custom Application Detectors 2629

Custom Application Detector and User-Defined Application Fields 2629

Configuring Custom Application Detectors 2632

Create a User-Defined Application 2633

Specifying Detection Patterns in Basic Detectors 2634

Specifying Detection Criteria in Advanced Detectors 2635

Specifying EVE Process Assignments 2636

Testing a Custom Application Protocol Detector 2637

Viewing or Downloading Detector Details 2638

Sorting the Detector List 2638

Filtering the Detector List 2638

Filter Groups for the Detector List 2639

Navigating to Other Detector Pages 2640

Activating and Deactivating Detectors 2640

Editing Custom Application Detectors 2641

Deleting Detectors 2642

CHAPTER 85 **Network Discovery Policies** 2643

Overview: Network Discovery Policies 2643

Requirements and Prerequisites for Network Discovery Policies 2644

Network Discovery Customization 2644

Configuring the Network Discovery Policy 2645

Network Discovery Rules 2645

Configuring Network Discovery Rules 2646

Actions and Discovered Assets 2646
Monitored Networks 2647
Port Exclusions 2649
Zones in Network Discovery Rules 2651
The Traffic-Based Detection Identity Source 2651
Configuring Advanced Network Discovery Options 2654
Network Discovery General Settings 2655
Configuring Network Discovery General Settings 2655
Network Discovery Identity Conflict Settings 2655
Configuring Network Discovery Identity Conflict Resolution 2656
Network Discovery Vulnerability Impact Assessment Options 2656
Enabling Network Discovery Vulnerability Impact Assessment 2657
Indications of Compromise 2657
Enabling Indications of Compromise Rules 2658
Adding NetFlow Exporters to a Network Discovery Policy 2658
Network Discovery Data Storage Settings 2659
Configuring Network Discovery Data Storage 2660
Configuring Network Discovery Event Logging 2661
Adding Network Discovery OS and Server Identity Sources 2661
Troubleshooting Your Network Discovery Strategy 2662

PART XIV **FlexConfig Policies** 2665

CHAPTER 86 **FlexConfig Policies** 2667
FlexConfig Policy Overview 2667
Recommended Usage for FlexConfig Policies 2668
CLI Commands in FlexConfig Objects 2668
Determine the ASA Software Version and Current CLI Configuration 2669
Prohibited CLI Commands 2669
Template Scripts 2671
FlexConfig Variables 2672
How to Process Variables 2672
How to See What a Variable Will Return for a Device 2675
FlexConfig Policy Object Variables 2676

FlexConfig System Variables 2677
Predefined FlexConfig Objects 2678
Predefined Text Objects 2683
Requirements and Prerequisites for FlexConfig Policies 2687
Guidelines and Limitations for FlexConfig 2687
Customizing Device Configuration with FlexConfig Policies 2688
Configure FlexConfig Objects 2690
Add a Policy Object Variable to a FlexConfig Object 2692
Configure Secret Keys 2693
Configure FlexConfig Text Objects 2694
Configure the FlexConfig Policy 2695
Set Target Devices for a FlexConfig Policy 2696
Preview the FlexConfig Policy 2697
Verify the Deployed Configuration 2698
Remove Features Configured Using FlexConfig 2699
Convert from FlexConfig to Managed Feature 2701
Examples for FlexConfig 2701
How to Configure Precision Time Protocol (ISA 3000) 2702
How to Configure Automatic Hardware Bypass for Power Failure (ISA 3000) 2706
Migrating FlexConfig Policies 2708
History for FlexConfig 2710

PART XV **Advanced Network Analysis in Snort 3** 2713

CHAPTER 87 **Advanced Access Control Settings for Network Analysis and Intrusion Policies** 2715
About Advanced Access Control Settings for Network Analysis and Intrusion Policies 2715
Requirements and Prerequisites for Advanced Access Control Settings for Network Analysis and Intrusion Policies 2715
Inspection of Packets That Pass Before Traffic Is Identified 2716
Best Practices for Handling Packets That Pass Before Traffic Identification 2716
Specify a Policy to Handle Packets That Pass Before Traffic Identification 2717
Advanced Settings for Network Analysis Policies 2717
Setting the Default Network Analysis Policy 2718
Network Analysis Rules 2719

Network Analysis Policy Rule Conditions 2719
Configuring Network Analysis Rules 2721
Managing Network Analysis Rules 2721

CHAPTER 88 **Transport and Network Layer Preprocessors** 2723
Introduction to Transport and Network Layer Preprocessors 2723
License Requirements for Transport and Network Layer Preprocessors 2723
Requirements and Prerequisites for Transport and Network Layer Preprocessors 2724
Advanced Transport/Network Preprocessor Settings 2724
Ignored VLAN Headers 2724
Active Responses in Intrusion Drop Rules 2724
Advanced Transport/Network Preprocessor Options 2725
Configuring Advanced Transport/Network Preprocessor Settings 2726

CHAPTER 89 **Adaptive Profiles** 2727
About Adaptive Profiles 2727
License Requirements for Adaptive Profiles 2728
Requirements and Prerequisites for Adaptive Profiles 2728
Adaptive Profile Updates 2728
Adaptive Profile Updates and Cisco Recommended Rules 2729
Adaptive Profile Options 2729
Configuring Adaptive Profiles 2730

PART XVI **Threat Intelligence Director** 2733

CHAPTER 90 **Secure Firewall Threat Intelligence Director** 2735
Secure Firewall Threat Intelligence Director Overview 2735
Threat Intelligence Director and Security Intelligence 2737
Performance Impact of Threat Intelligence Director 2737
Requirements and Prerequisites for Threat Intelligence Director 2738
Platform, Element, and License Requirements 2738
Source Requirements 2739
Source Content Limitations 2740
How To Set Up Threat Intelligence Director 2740

Configure Policies to Support Threat Intelligence Director 2741
Options for Ingesting Data Sources 2742
Fetch TAXII Feeds to Use as Sources 2743
Fetch Sources from a URL 2744
Upload a Local File to Use as a Source 2745
Handling of Duplicate Indicators 2747
Configure TLS/SSL Settings for a Threat Intelligence Director Source 2747
User Roles with Threat Intelligence Director Access 2748
About Backing Up and Restoring Threat Intelligence Director Data 2748
Analyze Threat Intelligence Director Incident and Observation Data 2749
Observation and Incident Generation 2749
View and Manage Incidents 2751
Incident Summary Information 2752
Incident Details 2752
View Events for a Threat Intelligence Director Observation 2756
Threat Intelligence Director Observations in Secure Firewall Management Center Events 2756
Factors That Affect the Action Taken 2757
Threat Intelligence Director-Firewall Management Center Action Prioritization 2757
View and Change Threat Intelligence Director Configurations 2761
View Threat Intelligence Director Status of Elements (Managed Devices) 2761
View and Manage Sources 2762
Source Summary Information 2763
Source Status Details 2764
View and Manage Indicators 2765
Indicator Summary Information 2766
Indicator Details 2767
View and Manage Observables 2768
Observable Summary Information 2769
Filter Threat Intelligence Director Data in Table Views 2769
Inheritance in Threat Intelligence Director Configurations 2770
Inheritance of TID Settings from Multiple Parents 2770
About Overriding Inherited TID Settings 2771
Edit Threat Intelligence Director Actions at the Source, Indicator, or Observable Level 2772
About Pausing Publishing 2772

Pause Threat Intelligence Director and Purge Threat Intelligence Director Data from Elements 2773
Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level 2774
Modify the Observable Publication Frequency 2775
About Adding Threat Intelligence Director Observables to the Do Not Block List 2775
Add Threat Intelligence Director Observables to a Do Not Block List 2776
View a STIX Source File 2776
Troubleshoot Threat Intelligence Director 2776
History for Threat Intelligence Director 2778

PART I

Getting Started with Device Configuration

- Device Management, on page 1
- Device Management Using Device Templates, on page 89
- Device Settings, on page 127
- Users, on page 213
- Change Management, on page 233
- Configuration Deployment, on page 247

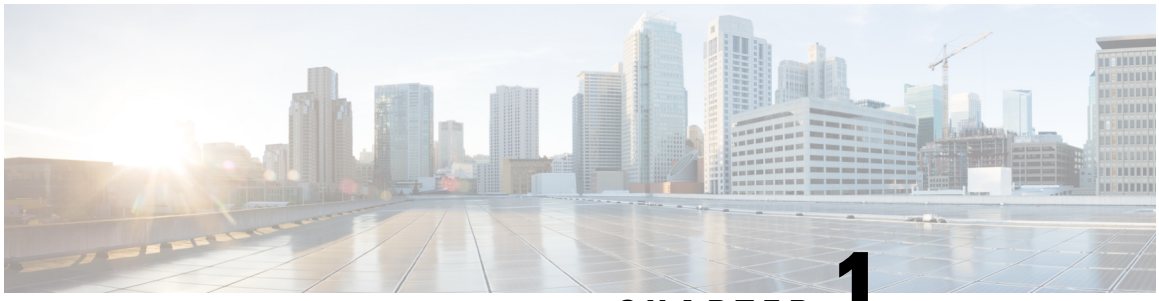

CHAPTER 1

Device Management

This guide applies to an *on-premises* Secure Firewall Management Center, either as your primary manager or as an analytics-only manager. When using the Cisco Security Cloud Control (Security Cloud Control) Cloud-Delivered Firewall Management Center as your primary manager, you can use an on-prem Firewall Management Center for analytics. Do not use this guide for cloud-delivered Firewall Management Center management; see Managing Firewall Threat Defense with Cloud-Delivered Firewall Management Center in Cisco Security Cloud Control.

You can add and manage devices in the Secure Firewall Management Center.

- About Device Management, on page 1
- Requirements and Prerequisites for Device Management, on page 10
- Log Into the Command-Line Interface on the Device, on page 11
- Complete the Firewall Threat Defense Initial Configuration for Manual Registration, on page 13
- Manage Devices, on page 28
- Switch Managers, on page 72
- Hot Swap an SSD on the Secure Firewall 3100/4200, on page 78
- Disable the USB Port, on page 81
- History for Device Management, on page 84

About Device Management

Use the Firewall Management Center to manage your devices.

About the Firewall Management Center and Device Management

When the Firewall Management Center manages a device, it sets up a two-way, SSL-encrypted communication channel between itself and the device. The Firewall Management Center uses this channel to send information to the device about how you want to analyze and manage your network traffic to the device. As the device evaluates the traffic, it generates events and sends them to the Firewall Management Center using the same channel.

By using the Firewall Management Center to manage devices, you can:

- configure policies for all your devices from a single location, making it easier to change configurations
- install various types of software updates on devices

- push health policies to your managed devices and monitor their health status from the Firewall Management Center

Note If you have a Security Cloud Control-managed device and are using the on-prem Firewall Management Center for analytics only, then the on-prem Firewall Management Center does not support policy configuration or upgrading. Chapters and procedures in this guide related to device configuration and other unsupported features do not apply to devices whose primary manager is Security Cloud Control.

The Firewall Management Center aggregates and correlates intrusion events, network discovery information, and device performance data, allowing you to monitor the information that your devices are reporting in relation to one another, and to assess the overall activity occurring on your network.

You can use the Firewall Management Center to manage nearly every aspect of a device's behavior.

Note Although the Firewall Management Center can manage devices running certain previous releases as specified in the compatibility matrix available at http://www.cisco.com/c/en/us/support/security/defense-center/products-device-support-tables-list.html, new features that require the latest version of Firewall Threat Defense software are not available to these previous-release devices. Some Firewall Management Center features may be available for earlier versions.

What Can Be Managed by a Secure Firewall Management Center?

You can use the Secure Firewall Management Center as a central management point to manage Firewall Threat Defense devices.

When you manage a device, information is transmitted between the Firewall Management Center and the device over a secure, TLS-1.3-encrypted communication channel. You do not need to run this traffic over an additional encrypted tunnel such as Site-to-Site VPN for security purposes. If the VPN goes down, for example, you will lose your management connection, so we recommend a simple management path.

The following illustration lists what is transmitted between the Firewall Management Center and its managed devices. Note that the types of events and policies that are sent between the appliances are based on the device type.

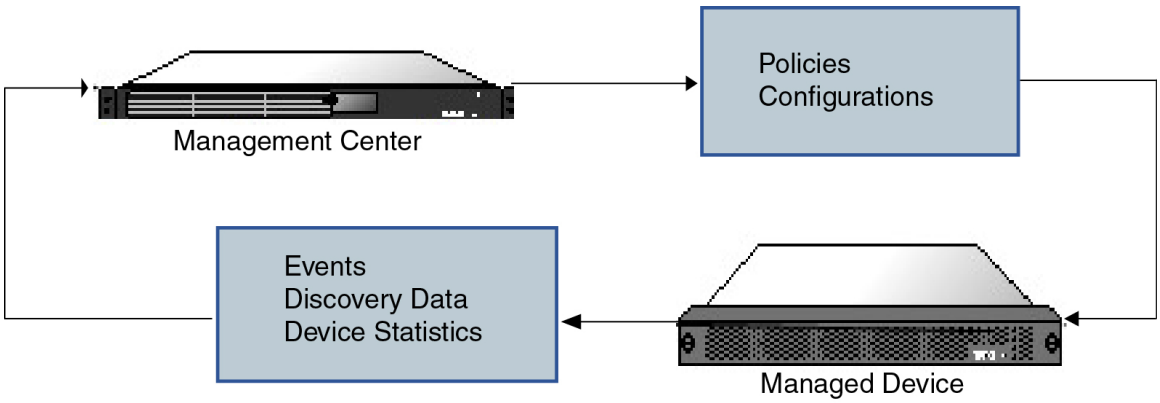

About the Management Connection

After you configure the device with the Firewall Management Center information and after you add the device to the Firewall Management Center, either the device or the Firewall Management Center can establish the management connection. Depending on initial setup:

- Either the device or the Firewall Management Center can initiate.
- Only the device can initiate.
- Only the Firewall Management Center can initiate.

Initiation always originates with eth0 on the Firewall Management Center or with the lowest-numbered management interface on the device. Additional management interfaces are tried if the connection is not established. Multiple management interfaces on the Firewall Management Center let you connect to discrete networks or to segregate management and event traffic. However, the initiator does not choose the best interface based on the routing table.

Make sure the management connection is stable, without excessive packet loss, with at least 5 Mbps throughput. By default, the management connection uses TCP port 8305 (this port is configurable). If you place another Firewall Threat Defense between devices and the Firewall Management Center, to prevent potential management disruption, be sure to exempt management traffic from deep inspection by applying a prefilter policy for it.

Note The management connection is a secure, TLS-1.3-encrypted communication channel between itself and the device. You do not need to run this traffic over an additional encrypted tunnel such as Site-to-Site VPN for security purposes. If the VPN goes down, for example, you will lose your management connection, so we recommend a simple management path.

Beyond Policies and Events

In addition to deploying policies to devices and receiving events from them, you can also perform other device-related tasks on the Firewall Management Center.

Backing Up a Device

You cannot backup a physical managed device from the FTD CLI. To back up configuration data, and, optionally, unified files, perform a backup of the device using the Firewall Management Center that is managing the device.

To back up event data, perform a backup of the Firewall Management Center that is managing the device.

Updating Devices

From time to time, Cisco releases updates to the Firepower System, including:

- intrusion rule updates, which may contain new and updated intrusion rules
- vulnerability database (VDB) updates
- geolocation updates
- software patches and updates

You can use the Firewall Management Center to install an update on the devices it manages.

About Device Management Interfaces

Each device includes a single dedicated Management interface for communicating with the Firewall Management Center. You can optionally configure the device to use a data interface for management instead of the dedicated Management interface.

You can perform initial setup on the management interface, or on the console port.

Management interfaces are also used to communicate with the Smart Licensing server, to download updates, and to perform other management functions.

Management and Event Interfaces on the Firewall Threat Defense

When you set up your device, you specify the Firewall Management Center IP address or hostname that you want to connect to, if known. In this case, the device initiates the connection, and both management and event traffic go to this address at initial registration. If the Firewall Management Center is not known, then the Firewall Management Center establishes the initial connection. In this case, it might initially connect from a different Firewall Management Center management interface than specified on the Firewall Threat Defense. Subsequent connections should use the Firewall Management Center management interface with the specified IP address.

If the Firewall Management Center has a separate event-only interface, the managed device sends subsequent event traffic to the Firewall Management Center event-only interface if the network allows. In addition, some managed-device models include an additional management interface that you can configure for event-only traffic. Note that if you configure a data interface for management, you cannot use separate management and event interfaces. If the event network goes down, then event traffic reverts to the regular management interfaces on the Firewall Management Center and/or on the managed device.

Using the Firewall Threat Defense Data Interface for Management

You can use either the dedicated Management interface or a regular data interface for communication with the Firewall Management Center. Manager access on a data interface is useful if you want to manage the Firewall Threat Defense remotely from the outside interface, or you do not have a separate management network. Moreover, using a data interface lets you configure a redundant secondary interface to take over management functions if the primary interface goes down.

Manager Access Requirements

Manager access from a data interface has the following requirements.

- You can only enable manager access on a physical, data interface. You cannot use a subinterface or EtherChannel, nor can you create a subinterface on the manager access interface. You can also use the Firewall Management Center to enable manager access on a single secondary interface for redundancy.
- This interface cannot be management-only.
- Routed firewall mode only, using a routed interface.
- PPPoE is not supported. If your ISP requires PPPoE, you will have to put a router with PPPoE support between the Firewall Threat Defense and the WAN modem.
- The interface must be in the global VRF only.

- SSH is not enabled by default for data interfaces, so you will have to enable SSH later using the Firewall Management Center. Because the Management interface gateway will be changed to be the data interfaces, you also cannot SSH to the Management interface from a remote network unless you add a static route for the Management interface using the **configure network static-routes** command. For Firewall Threat Defense Virtual on Amazon Web Services, a console port is not available, so you should maintain your SSH access to the Management interface: add a static route for Management before you continue with your configuration. Alternatively, be sure to finish all CLI configuration (including the **configure manager add** command) before you configure the data interface for manager access and you are disconnected.
- You cannot use separate management and event-only interfaces.
- Clustering is not supported. You must use the Management interface in this case.

High Availability Requirements

When using a data interface with device high availability, see the following requirements.

- Use the same data interface on both devices for manager access.
- You cannot use DHCP; only a static IP address is supported. Features that rely on DHCP cannot be used, including DDNS and zero-touch provisioning.

Note If you use zero-touch provisioning to register the device, when you use the outside interface for manager access, it uses DHCP by default. Before you can enable high availability, you need to change the IP address to a static address. See Change the Device IP Address, on page 158. Alternatively, you can use the Management interface instead; DHCP is supported on Management with high availability.

- Have different static IP addresses in the same subnet.
- Use the same manager configuration (**configure manager add** command) to ensure that the connectivity is the same.
- You cannot use the data interface as the failover or state link.

Management Interface Support Per Device Model

See the hardware installation guide for your model for the management interface locations.

Note For the Firepower 4100/9300, the MGMT interface is for *chassis* management, not for Firewall Threat Defense logical device management. You must configure a separate interface to be of type mgmt (and/or firepower-eventing), and then assign it to the Firewall Threat Defense logical device.

See the following table for supported management interfaces on each managed device model.

Table 1: Management Interface Support on Managed Devices

Model	Management Interface	Optional Event Interface
Firepower 1000	management0 **Note** management0 is the internal name of the Management 1/1 interface.	No Support
Secure Firewall 1200	management0 **Note** management0 is the internal name of the Management 1/1 interface.	No Support
Secure Firewall 3100	management0 **Note** management0 is the internal name of the Management 1/1 interface.	No Support
Secure Firewall 4200	management0 **Note** management0 is the internal name of the Management 1/1 interface.	management1 **Note** management1 is the internal name of the Management 1/2 interface.
Firepower 4100 and 9300	management0 **Note** management0 is the internal name of this interface, regardless of the physical interface ID.	management1 **Note** management1 is the internal name of this interface, regardless of the physical interface ID.
ISA 3000	br1 **Note** br1 is the internal name of the Management 1/1 interface.	No support
Secure Firewall Threat Defense Virtual	eth0	No support

Network Routes on Device Management Interfaces

Management interfaces (including event-only interfaces) support only static routes to reach remote networks. When you set up your managed device, the setup process creates a default route to the gateway IP address that you specify. You cannot delete this route; you can only modify the gateway address.

Note The routing for management interfaces is completely separate from routing that you configure for data interfaces. If you configure a data interface for management instead of using the dedicated Management interface, traffic is routed over the backplane to use the data routing table. The information in this section does not apply.

You can configure multiple management interfaces on some platforms (a management interface and an event-only interface). The default route does not include an egress interface, so the interface chosen depends on the gateway address you specify, and which interface's network the gateway belongs to. In the case of multiple interfaces on the default network, the device uses the lower-numbered interface as the egress interface.

At least one static route is recommended per management interface to access remote networks. We recommend placing each interface on a separate network to avoid potential routing problems, including routing problems from other devices to the Firewall Threat Defense.

Note The interface used for management connections is not determined by the routing table. Connections are always tried using the lowest-numbered interface first.

NAT Environments

Network address translation (NAT) is a method of transmitting and receiving network traffic through a router that involves reassigning the source or destination IP address. The most common use for NAT is to allow private networks to communicate with the internet. Static NAT performs a 1:1 translation, which does not pose a problem for Firewall Management Center communication with devices, but port address translation (PAT) is more common. PAT lets you use a single public IP address and unique ports to access the public network; these ports are dynamically assigned as needed, so you cannot initiate a connection to a device behind a PAT router.

Normally, you need both IP addresses (along with a registration key) for both routing purposes and for authentication: the Firewall Management Center specifies the device IP address when you add a device, and the device specifies the Firewall Management Center IP address. However, if you only know one of the IP addresses, which is the minimum requirement for routing purposes, then you must also specify a unique NAT ID on both sides of the connection to establish trust for the initial communication and to look up the correct registration key. The Firewall Management Center and device use the registration key and NAT ID (instead of IP addresses) to authenticate and authorize for initial registration.

For example, you add a device to the Firewall Management Center, and you do not know the device IP address (for example, the device is behind a PAT router), so you specify only the NAT ID and the registration key on the Firewall Management Center; leave the IP address blank. On the device, you specify the Firewall Management Center IP address, the same NAT ID, and the same registration key. The device registers to the Firewall Management Center's IP address. At this point, the Firewall Management Center uses the NAT ID instead of IP address to authenticate the device.

Although the use of a NAT ID is most common for NAT environments, you might choose to use the NAT ID to simplify adding many devices to the Firewall Management Center. On the Firewall Management Center, specify a unique NAT ID for each device you want to add while leaving the IP address blank, and then on each device, specify both the Firewall Management Center IP address and the NAT ID. Note: The NAT ID must be unique per device.

The following example shows three devices behind a PAT IP address. In this case, specify a unique NAT ID per device on both the Firewall Management Center and the devices, and specify the Firewall Management Center IP address on the devices.

Figure 1: NAT ID for Managed Devices Behind PAT

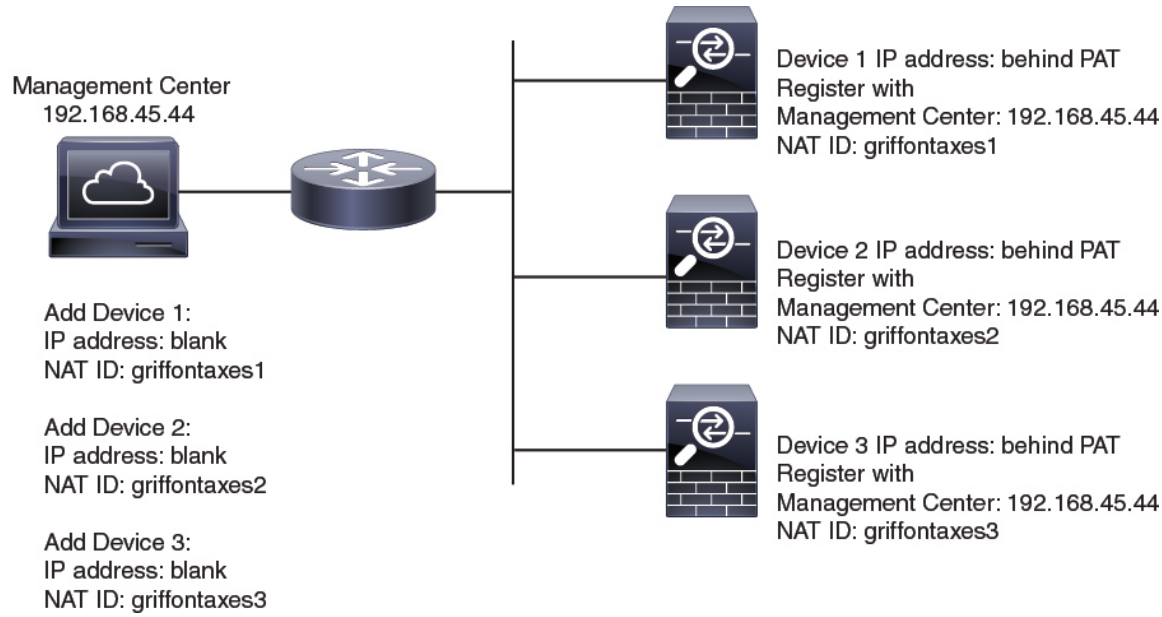

The following example shows the Firewall Management Center behind a PAT IP address. In this case, specify a unique NAT ID per device on both the Firewall Management Center and the devices, and specify the device IP addresses on the Firewall Management Center.

Figure 2: NAT ID for Firewall Management Center Behind PAT

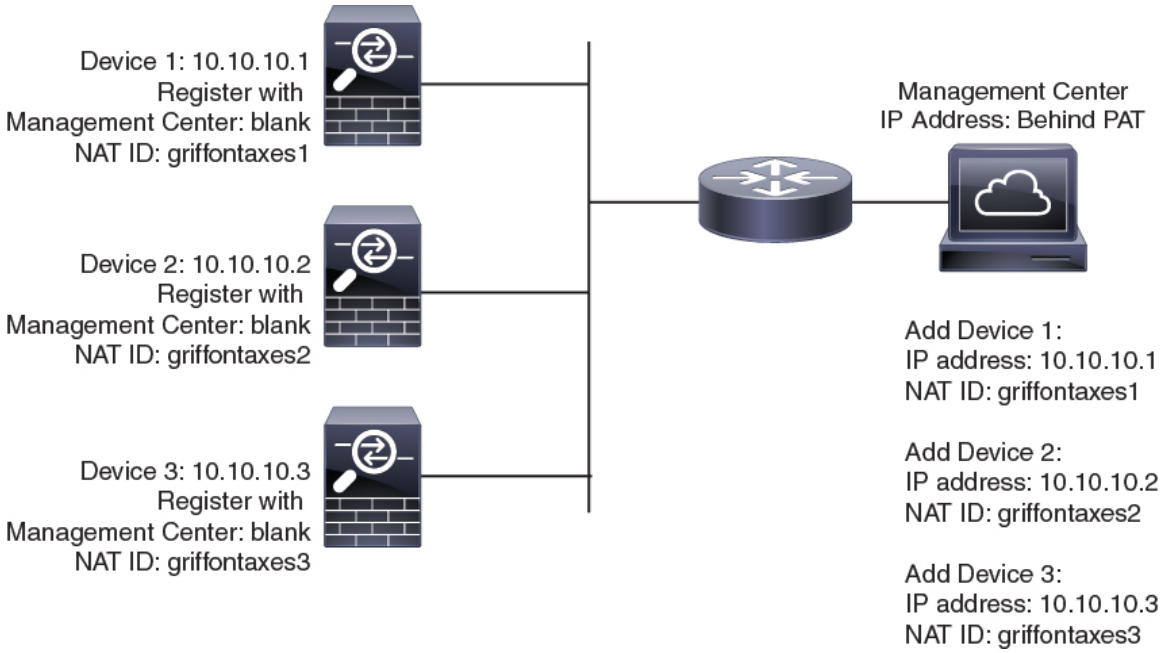

Management and Event Traffic Channel Examples

Note If you use a data interface for management on a Firewall Threat Defense, you cannot use separate management and event interfaces for that device.

The following example shows the Firewall Management Center and managed devices using only the default management interfaces.

Figure 3: Single Management Interface on the Secure Firewall Management Center

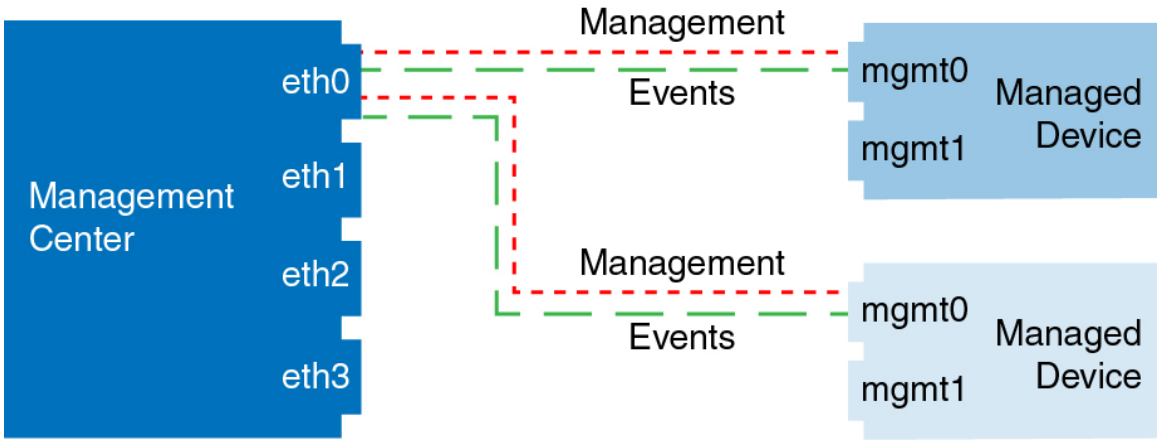

The following example shows the Firewall Management Center using separate management interfaces for devices; and each managed device using 1 management interface.

Figure 4: Multiple Management Interfaces on the Secure Firewall Management Center

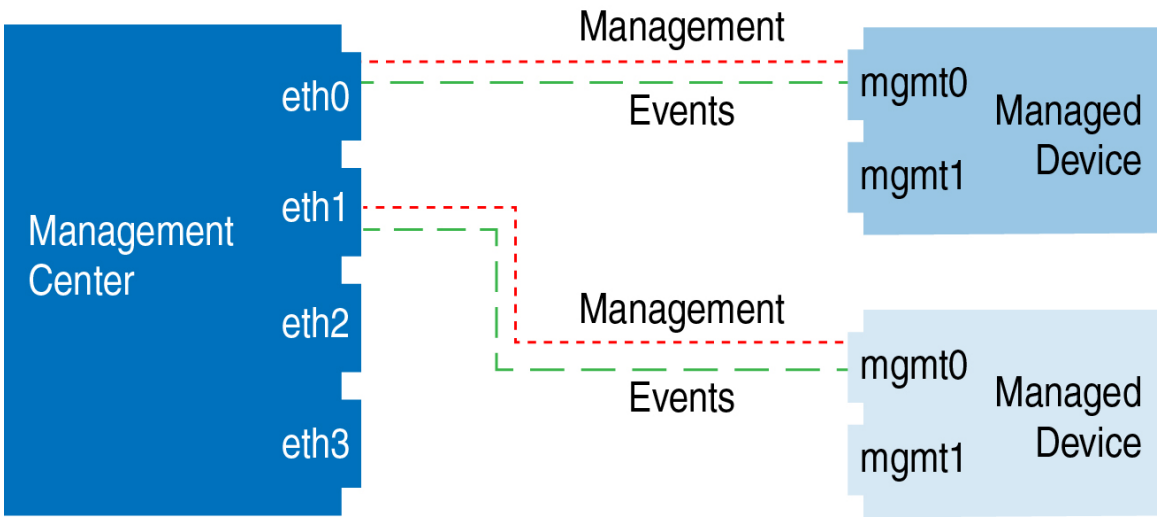

The following example shows the Firewall Management Center and managed devices using a separate event interface.

Figure 5: Separate Event Interface on the Secure Firewall Management Center and Managed Devices

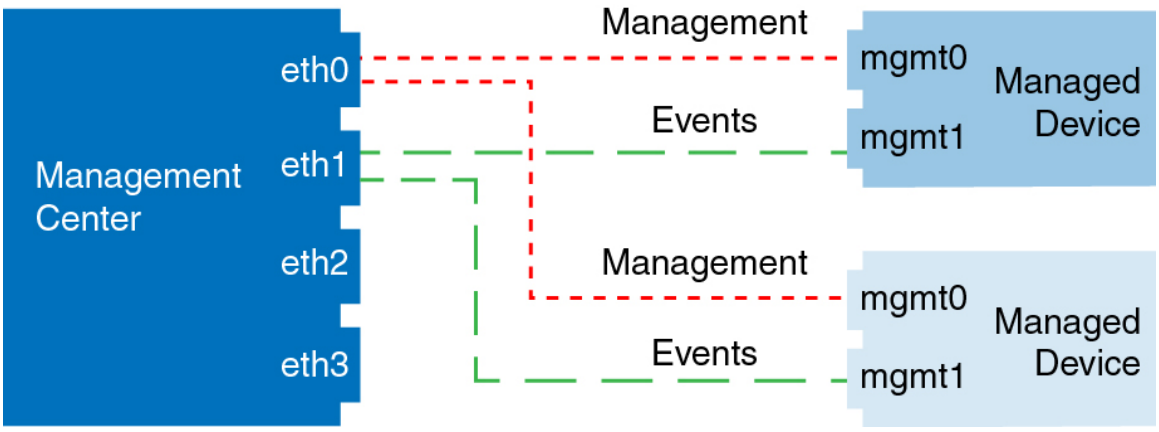

The following example shows a mix of multiple management interfaces and a separate event interface on the Firewall Management Center and a mix of managed devices using a separate event interface, or using a single management interface.

Figure 6: Mixed Management and Event Interface Usage

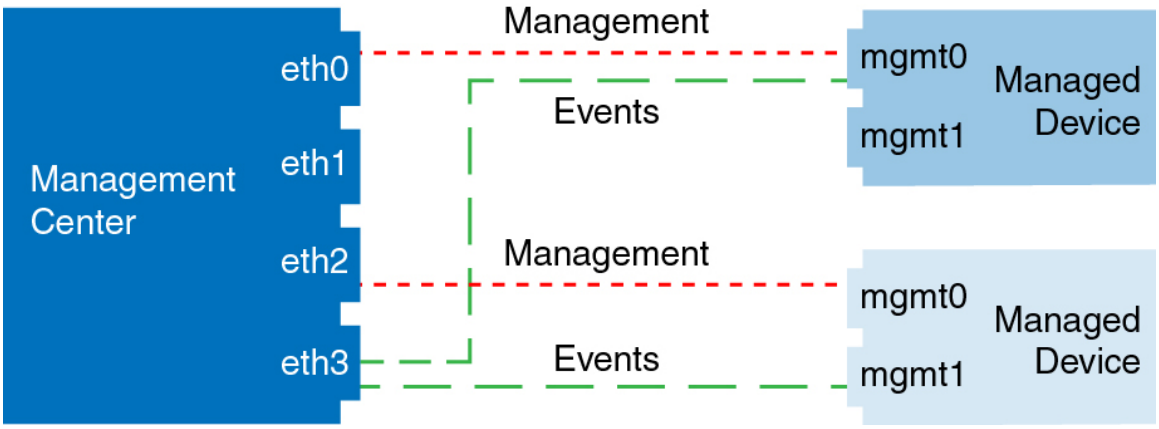

Requirements and Prerequisites for Device Management

Supported Domains

The domain in which the device resides.

User Roles

- Admin
- Network Admin

Management Connection

Make sure the management connection is stable, without excessive packet loss, with at least 5Mbps throughput.

Zero-Touch Provisioning Requirements

Zero-Touch Provisioning is not supported with clustering or multi-instance mode.

High availability is only supported when you use the Management interface because zero-touch provisioning uses DHCP, which is not supported for data interfaces and high availability.

Zero-Touch Provisioning is supported on the following models using 7.4 or later:

- Firepower 1010
- Firepower 1100
- Secure Firewall 1200
- Firepower 2100 (on supported device versions)
- Secure Firewall 3100

Log Into the Command-Line Interface on the Device

You can log directly into the command-line interface on Firewall Threat Defense devices. If this is your first time logging in, complete the initial setup process using the default **admin** user; see Complete the Firewall Threat Defense Initial Configuration Using the CLI, on page 19.

Note If a user makes three consecutive failed attempts to log into the CLI via SSH, the system terminates the SSH connection.

Before you begin

- Create additional user accounts that can log into the CLI using the **configure user add** command.
- If you get unreadable characters when connecting to the console port, verify the port settings. If they are correct, try the cable with another device using the same settings. If the cable is good, you might need to replace the hardware for the console port. Also consider trying a different workstation to make the connection.

Procedure

Step 1 Connect to the Firewall Threat Defense CLI, either from the console port or using SSH.

You can SSH to the management interface of the Firewall Threat Defense device. You can also connect to the address on a data interface if you open the interface for SSH connections. SSH access to data interfaces is disabled by default. See SSH Access, on page 964 to allow SSH connections to specific data interfaces.

For physical devices, you can directly connect to the console port on the device. See the hardware guide for your device for more information about the console cable. Use the following serial settings:

- 9600 baud
- 8 data bits

- No parity
- 1 stop bit

The CLI on the console port is FXOS (with the exception of the ISA 3000, where it is the regular Firewall Threat Defense CLI). Use the Firewall Threat Defense CLI for basic configuration, monitoring, and normal system troubleshooting. See the FXOS documentation for information on FXOS commands.

For a chassis in multi-instance mode, you can connect to FXOS on the console port, or you can enable SSH for the Management interface according to Configure SSH and SSH Access List, on page 397. SSH is disabled by default.

Step 2 Log in with the **admin** username and password.

Example:

```
firepower login: admin
Password:
Last login: Thu May 16 14:01:03 UTC 2019 on ttyS0
Successful login attempts for user 'admin' : 1

firepower#
```

Step 3 If you used the console port, access the Firewall Threat Defense CLI.

connect ftd

Multi-instance mode:

connect ftd *name*

To view the instance names, enter the command without a name.

Note
This step does not apply to the ISA 3000.

Example:

```
firepower# connect ftd
>
```

Step 4 At the CLI prompt (>), use any of the commands allowed by your level of command line access.

To return to FXOS on the console port, enter **exit**.

Step 5 (Optional) If you used SSH, you can connect to FXOS.

connect fxos

To return to the Firewall Threat Defense CLI, enter **exit**.

Step 6 (Optional) Access the diagnostic CLI:

system support diagnostic-cli

Use this CLI for advanced troubleshooting. This CLI includes additional **show** and other commands.

This CLI has submodes: user EXEC mode, privileged EXEC mode, and recovery-config mode. More commands are available in privileged EXEC mode than user EXEC mode. To enter privileged EXEC mode, enter the **enable** command; press enter without entering a password when prompted.

Example:

```
> system support diagnostic-cli
firepower> enable
Password:
firepower#
```

To use recovery-config mode, see Access Recovery-Config Mode in the Diagnostic CLI, on page 146.

To return to the regular CLI, type **Ctrl-a**, **d**.

Complete the Firewall Threat Defense Initial Configuration for Manual Registration

You can complete the Firewall Threat Defense initial configuration using the CLI or the Firewall Device Manager for all models except for the Firepower 4100/9300. For the Firepower 4100/9300, you complete initial configuration when you deploy the logical device. See Logical Devices on the Firepower 4100/9300, on page 297.

For zero-touch provisioning (serial number registration), you should not log into the device or perform initial setup. See Add a Device Using the Serial Number (Zero-Touch Provisioning)—Basic Configuration, on page 41.

Complete the Firewall Threat Defense Initial Configuration Using the Firewall Device Manager

When you use the Firewall Device Manager for initial setup, the following interfaces are preconfigured in addition to the Management interface and manager access settings:

- Ethernet 1/1—"outside", IP address from DHCP, IPv6 autoconfiguration
- Ethernet 1/2 (or for the Firepower 1010 and Secure Firewall 1210/1220, the VLAN1 interface)— "inside", 192.168.95.1/24
- Default route—Obtained through DHCP on the outside interface

Note that other settings, such as the DHCP server on inside, access control policy, or security zones, are not configured.

If you perform additional interface-specific configuration within Firewall Device Manager before registering with the Firewall Management Center, then that configuration is preserved.

When you use the CLI, only the Management interface and manager access settings are retained (for example, the default inside interface configuration is not retained).

- The Secure Firewall 4200 does not support the Firewall Device Manager. You need to use the CLI procedure: Complete the Firewall Threat Defense Initial Configuration Using the CLI, on page 19.
- This procedure does not apply for Security Cloud Control-managed devices for which you want to use an on-prem Firewall Management Center *for analytics only*. The Firewall Device Manager configuration

is meant to configure the primary manager. See Complete the Firewall Threat Defense Initial Configuration Using the CLI, on page 19 for more information about configuring the device for analytics.

- This procedure applies to all other devices except for the Firepower 4100/9300 and the ISA 3000. You can use the Firewall Device Manager to onboard these devices to the Firewall Management Center, but because they have different default configurations than other platforms, the details in this procedure may not apply to these platforms.

Procedure

Step 1 Log into the Firewall Device Manager.

a) Enter the following URL in your browser.

- Inside—**https://192.168.95.1**.
- Management—**https://***management_ip*. The Management interface is a DHCP client, so the IP address depends on your DHCP server. You will have to set the Management IP address to a static address as part of this procedure, so we recommend that you use the inside interface so you do not become disconnected.

b) Log in with the username **admin**, and the default password **Admin123**.

c) You are prompted to read and accept the End User License Agreement and change the admin password.

Step 2 Use the setup wizard when you first log into the Firewall Device Manager to complete the initial configuration. You can optionally skip the setup wizard by clicking **Skip device setup** at the bottom of the page.

After you complete the setup wizard, in addition to the default configuration for the inside interface, you will have configuration for an outside (Ethernet1/1) interface that will be maintained when you switch to the Firewall Management Center management.

a) Configure the following options for the outside and management interfaces, and click **Next**.

1. **Outside Interface Address**—This interface is typically the internet gateway, and might be used as your manager access interface. You cannot select an alternative outside interface during initial device setup. The first data interface is the default outside interface.

   If you want to use a different interface from outside (or inside) for manager access, you will have to configure it manually after completing the setup wizard.

   Configure IPv4—The IPv4 address for the outside interface. You can use DHCP or manually enter a static IP address, subnet mask, and gateway. You can also select **Off** to not configure an IPv4 address. You cannot configure PPPoE using the setup wizard. PPPoE may be required if the interface is connected to a DSL modem, cable modem, or other connection to your ISP, and your ISP uses PPPoE to provide your IP address. You can configure PPPoE after you complete the wizard.

   Configure IPv6—The IPv6 address for the outside interface. You can use DHCP or manually enter a static IP address, prefix, and gateway. You can also select **Off** to not configure an IPv6 address.

2. **Management Interface**

   You will not see Management Interface settings if you performed initial setup at the CLI.

   The Management interface settings are used even if you enable manager access on a data interface. For example, the management traffic that is routed over the backplane through the data interface will resolve FQDNs using the Management interface DNS servers, and not the data interface DNS servers.

DNS Servers—The DNS server for the system's management address. Enter one or more addresses of DNS servers for name resolution. The default is the OpenDNS public DNS servers. If you edit the fields and want to return to the default, click **Use OpenDNS** to reload the appropriate IP addresses into the fields.

Firewall Hostname—The hostname for the system's management address.

b) Configure the **Time Setting (NTP)** and click **Next**.

1. **Time Zone**—Select the time zone for the system.
2. **NTP Time Server**—Select whether to use the default NTP servers or to manually enter the addresses of your NTP servers. You can add multiple servers to provide backups.

c) Select **Start 90 day evaluation period without registration**.

Do not register the Firewall Threat Defense with the Smart Software Manager; all licensing is performed on the Firewall Management Center.

d) Click **Finish**.

e) You are prompted to choose **Cloud Management** or **Standalone**. For Firewall Management Center management, choose **Standalone**, and then **Got It**.

Step 3 (Might be required) Configure the Management interface.

You may need to change the Management interface configuration, even if you intend to use a data interface for manager access. You will have to reconnect to the Firewall Device Manager if you were using the Management interface for the Firewall Device Manager connection.

- Data interface for manager access—The Management interface must have the gateway set to data interfaces. By default, the Management interface receives an IP address and gateway from DHCP. If you do not receive a gateway from DHCP (for example, you did not connect this interface to a network), then the gateway will default to data interfaces, and you do not need to configure anything. If you did receive a gateway from DHCP, then you need to instead configure this interface with a static IP address and set the gateway to data interfaces.
- Management interface for manager access—If you want to configure a static IP address, be sure to also set the default gateway to be a unique gateway instead of the data interfaces. If you use DHCP, then you do not need to configure anything assuming you successfully get the gateway from DHCP.

Step 4 If you want to configure additional interfaces, including an interface other than outside or inside that you want to use for manager access, choose **Device**, and then click the link in the **Interfaces** summary.

Other Firewall Device Manager configuration will not be retained when you register the device to Firewall Management Center.

Step 5 Choose **Device** > **System Settings** > **Central Management**, and click **Proceed** to set up the Firewall Management Center management.

Step 6 Configure the **Management Center/SCC Details**.

Figure 7: Management Center/SCC Details

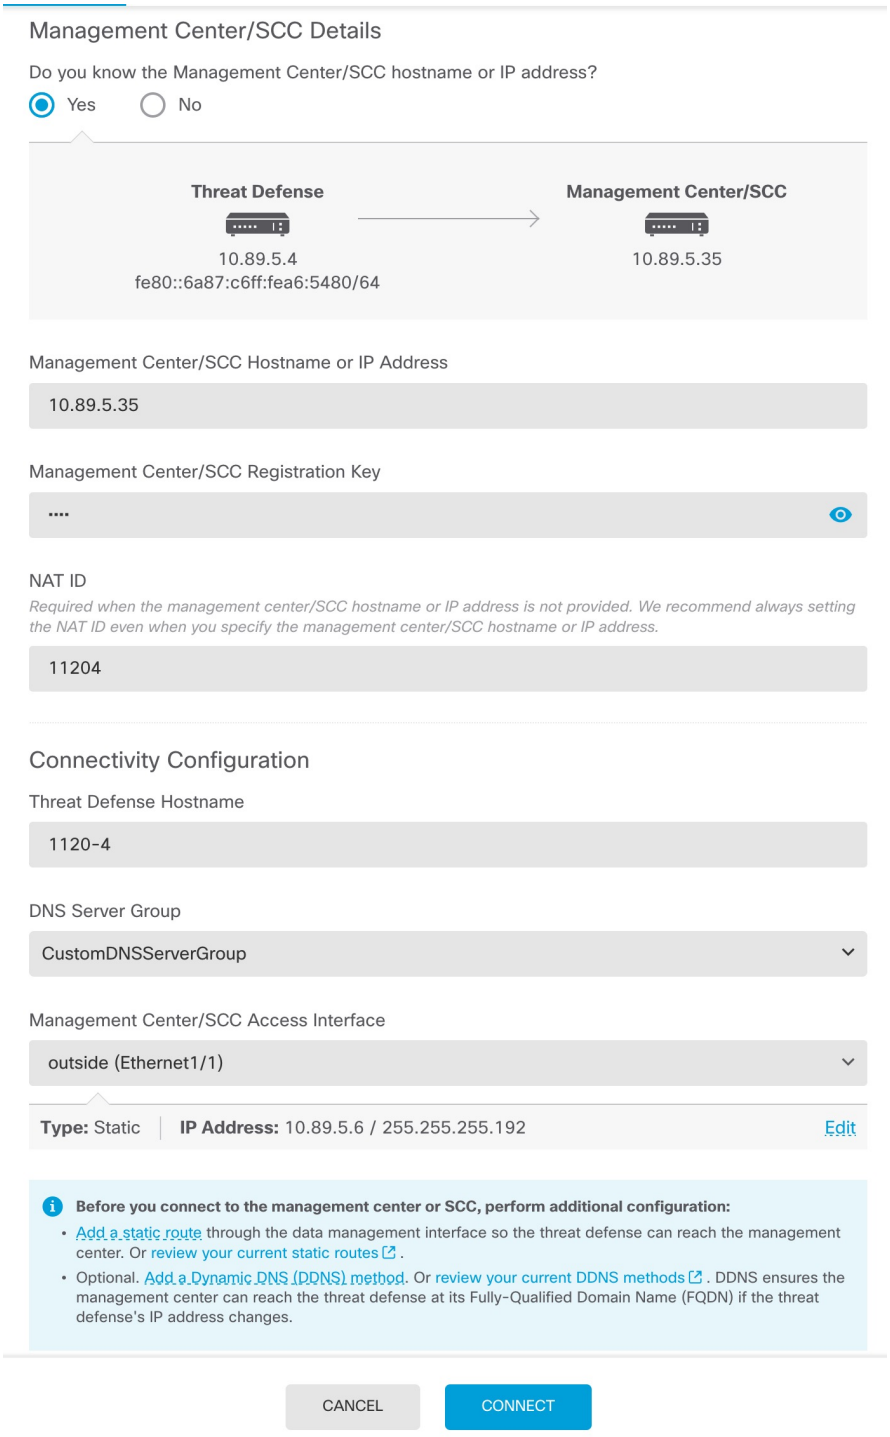

a) For **Do you know the Management Center/SCC hostname or IP address?**, click **Yes** if you can reach the Firewall Management Center using an IP address or hostname, or **No** if the Firewall Management CenterSecurity Cloud Control is behind NAT or does not have a public IP address or hostname.

At least one of the devices, either the Firewall Management Center or the Firewall Threat Defense device, must have a reachable IP address to establish the two-way, TLS-1.3-encrypted communication channel between the two devices.

b) If you chose **Yes**, then enter the **Management Center/SCC Hostname or IP Address**.

c) Specify the **Management Center/SCC Registration Key**.

This key is a one-time registration key of your choice that you will also specify on the Firewall Management Center when you register the Firewall Threat Defense device. The registration key must be between 2 and 36 characters. Valid characters include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-). This ID can be used for multiple devices registering to the Firewall Management Center.

a) Specify a **NAT ID**.

This ID is a unique, one-time string of your choice that you will also specify on the Firewall Management Center. The NAT ID must be between 2 and 36 characters. Valid characters include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-). This ID *cannot* be used for any other devices registering to the Firewall Management Center. The NAT ID is used in combination with the IP address to verify that the connection is coming from the correct device; only after authentication of the IP address/NAT ID will the registration key be checked. We recommended that you always use the NAT ID even when it is optional, but it is required if:

- You set the Firewall Management Center IP address to **DONTRESOLVE**.
- When adding the device on the Firewall Management Center, you do not specify a reachable device IP address or hostname.
- You use the data interface for management, even if you specify IP addresses on both sides.
- The Firewall Management Center uses multiple management interfaces.

Step 7 Configure the **Connectivity Configuration**.

a) Specify the **FTD Hostname**.

If you use a data interface for the **Management Center/SCC Access Interface** access, then this FQDN will be used for this interface.

b) Specify the **DNS Server Group**.

Choose an existing group, or create a new one. The default DNS group is called **CiscoUmbrellaDNSServerGroup**, which includes the OpenDNS servers.

If you intend to choose a data interface for the **Management Center/SCC Access Interface**, then this setting sets the *data* interface DNS server. The Management DNS server that you set with the setup wizard is used for management traffic. The data DNS server is used for DDNS (if configured) or for security policies applied to this interface. You are likely to choose the same DNS server group that you used for Management, because both management and data traffic reach the DNS server through the outside interface.

On the Firewall Management Center, the data interface DNS servers are configured in the Platform Settings policy that you assign to this Firewall Threat Defense device. When you add the Firewall Threat Defense device to the Firewall Management Center, the local setting is maintained, and the DNS servers are *not* added to a Platform Settings policy. However, if you later assign a Platform Settings policy to the Firewall Threat Defense device that includes a DNS configuration, then that configuration will overwrite the local setting. We suggest that you actively configure the DNS Platform Settings to match this setting to bring the Firewall Management Center and the Firewall Threat Defense device into sync.

Also, local DNS servers are only retained by the Firewall Management Center if the DNS servers were discovered at initial registration.

If you intend to choose the Management interface for the **Management Center/SCC Access InterfaceFMC Access Interface**, then this setting configures the Management DNS server.

c) For the **Management Center/SCC Access Interface**, choose any configured interface.

You can change the manager interface after you register the Firewall Threat Defense device to the Firewall Management Center, to either the Management interface or another data interface.

Step 8 (Optional) If you chose a data interface, and it was not the outside interface, then add a default route.

You will see a message telling you to check that you have a default route through the interface. If you chose outside, you already configured this route as part of the setup wizard. If you chose a different interface, then you need to manually configure a default route before you connect to the Firewall Management Center.

If you chose the Management interface, then you need to configure the gateway to be a unique gateway before you can proceed on this screen.

Step 9 (Optional) If you chose a data interface, click **Add a Dynamic DNS (DDNS) method**.

DDNS ensures the Firewall Management Center can reach the Firewall Threat Defense device at its Fully-Qualified Domain Name (FQDN) if the IP address changes. See **Device** > **System Settings** > **DDNS Service** to configure DDNS.

If you configure DDNS before you add the Firewall Threat Defense device to the Firewall Management Center, the Firewall Threat Defense device automatically adds certificates for all of the major CAs from the Cisco Trusted Root CA bundle so that the Firewall Threat Defense device can validate the DDNS server certificate for the HTTPS connection. Firewall Threat Defense supports any DDNS server that uses the DynDNS Remote API specification (https://help.dyn.com/remote-access-api/).

DDNS is not supported when using the Management interface for manager access.

Step 10 Click **Connect**. The **Registration Status** dialog box shows the current status of the switch to the Firewall Management Center. After the **Saving Management Center/SCC Registration Settings** step, go to the Firewall Management Center, and add the firewall.

If you want to cancel the switch to the Firewall Management Center, click **Cancel Registration**. Otherwise, do not close the Firewall Device Manager browser window until after the **Saving Management Center/SCC Registration Settings** step. If you do, the process will be paused, and will only resume when you reconnect to the Firewall Device Manager.

If you remain connected to the Firewall Device Manager after the **Saving Management Center/SCC Registration Settings** step, you will eventually see the **Successful Connection with Management Center/SCC** dialog box, after which you will be disconnected from the Firewall Device Manager.

Figure 8: Successful Connection

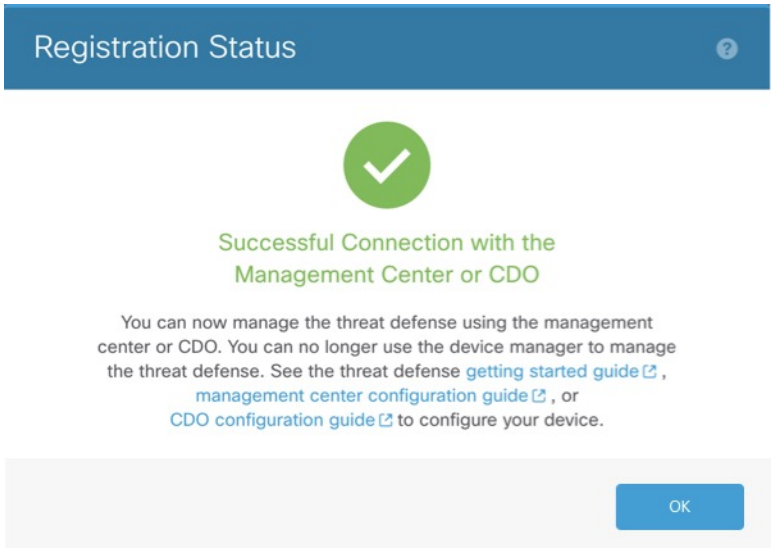

Complete the Firewall Threat Defense Initial Configuration Using the CLI

Connect to the Firewall Threat Defense CLI to perform initial setup, including setting the Management IP address, gateway, and other basic networking settings using the setup wizard. The dedicated Management interface is a special interface with its own network settings. If you do not want to use the Management interface for manager access, you can use the CLI to configure a data interface instead. You will also configure Firewall Management Center communication settings. When you perform initial setup using the Firewall Device Manager, *all* interface configuration completed in the Firewall Device Manager is retained when you switch to the Firewall Management Center for management, in addition to the Management interface and manager access interface settings. Note that other default configuration settings, such as the access control policy, are not retained.

This procedure applies to all models except for the Firepower 4100/9300. To deploy a logical device and complete initial configuration on the Firepower 4100/9300, see Logical Devices on the Firepower 4100/9300, on page 297.

Procedure

Step 1 Connect to the Firewall Threat Defense CLI, either from the console port or using SSH to the Management interface, which obtains an IP address from a DHCP server by default. If you intend to change the network settings, we recommend using the console port so you do not get disconnected.

(Firepower and Secure Firewall hardware models) The console port connects to the FXOS CLI. The SSH session connects directly to the Firewall Threat Defense CLI.

Step 2 Log in with the username **admin** and the password **Admin123**.

(Firepower and Secure Firewall hardware models) At the console port, you connect to the FXOS CLI. The first time you log in to FXOS, you are prompted to change the password. This password is also used for the Firewall Threat Defense login for SSH.

Note
If the password was already changed, and you do not know it, you must reimage the device to reset the password to the default.

For Firepower and Secure Firewall hardware, see the Reimage Procedures in the Cisco FXOS Troubleshooting Guide for the Firepower 1000/2100 and Secure Firewall 3100/4200 with Threat Defense .

For the ASA 5500-X and ISA 3000, see the Cisco Secure Firewall ASA and Secure Firewall Threat Defense Reimage Guide.

Example:

```
firepower login: admin
Password: Admin123
Successful login attempts for user 'admin' : 1

[...]

Hello admin. You must change your password.
Enter new password: ********
Confirm new password: ********
Your password was updated successfully.

[...]

firepower#
```

Step 3 (Firepower and Secure Firewall hardware models) If you connected to FXOS on the console port, connect to the Firewall Threat Defense CLI.

connect ftd

Example:

```
firepower# connect ftd
>
```

Step 4 The first time you log in to the Firewall Threat Defense, you are prompted to accept the End User License Agreement (EULA) and, if using an SSH connection, to change the admin password. You are then presented with the CLI setup script.

Note
You cannot repeat the CLI setup wizard unless you clear the configuration; for example, by reimaging. However, all of these settings can be changed later at the CLI using **configure network** commands. See the threat defense command reference.

Defaults or previously entered values appear in brackets. To accept previously entered values, press **Enter**.

Note
The Management interface settings are used even when you enable manager access on a data interface. For example, the management traffic that is routed over the backplane through the data interface will resolve FQDNs using the Management interface DNS servers, and not the data interface DNS servers.

See the following guidelines:

- **Do you want to configure IPv4?** and/or **Do you want to configure IPv6?**—Enter **y** for at least one of these types of addresses.
- **Enter the IPv4 default gateway for the management interface** and/or **Enter the IPv6 gateway for the management interface**—If you want to use a data interface for manager access instead of the Management interface, choose **manual**. Although you do not plan to use the Management interface, you must set an IP address, for example, a private address. Make sure this interface is on a different subnet from the manager access interface to prevent routing issues. You cannot configure a data interface for management if the management interface is set to DHCP, because the default route, which must be **data-interfaces** (see the next bullet), might be overwritten with one received from the DHCP server.
- **Enter the IPv4 default gateway for the management interface** and/or **Configure IPv6 via DHCP, router, or manually?**—If you want to use a data interface for manager access instead of the management interface, set the gateway to be **data-interfaces**. This setting forwards management traffic over the backplane so it can be routed through the manager access data interface. If you want to use the Management interface for manager access, you should set a gateway IP address on the Management 1/1 network.
- **If your networking information has changed, you will need to reconnect**—If you are connected with SSH but you change the IP address at initial setup, you will be disconnected. Reconnect with the new IP address and password. Console connections are not affected.
- **Manage the device locally?**—Enter **no** to use the Firewall Management Center. A **yes** answer means you will use Secure Firewall Device Manager instead.
- **Configure firewall mode?**—We recommend that you set the firewall mode at initial configuration. Changing the firewall mode after initial setup erases your running configuration. Note that data interface manager access is only supported in routed firewall mode.

Example:

```
You must accept the EULA to continue.
Press <ENTER> to display the EULA:
Cisco General Terms
[...]

Please enter 'YES' or press <ENTER> to AGREE to the EULA:


System initialization in progress.  Please stand by.
You must configure the network to continue.
Configure at least one of IPv4 or IPv6 unless managing via data interfaces.
Do you want to configure IPv4? (y/n) [y]:
Do you want to configure IPv6? (y/n) [y]: n
Configure IPv4 via DHCP or manually? (dhcp/manual) [manual]:
Enter an IPv4 address for the management interface [192.168.45.61]: 10.89.5.17
Enter an IPv4 netmask for the management interface [255.255.255.0]: 255.255.255.192
Enter the IPv4 default gateway for the management interface [data-interfaces]: 10.89.5.1
Enter a fully qualified hostname for this system [firepower]: 1010-3
Enter a comma-separated list of DNS servers or 'none'
[208.67.222.222,208.67.220.220,2620:119:35::35]:
Enter a comma-separated list of search domains or 'none' []: cisco.com
If your networking information has changed, you will need to reconnect.
Disabling IPv6 configuration: management0
Setting DNS servers: 208.67.222.222,208.67.220.220,2620:119:35::35
Setting DNS domains:cisco.com
Setting hostname as 1010-3
Setting static IPv4: 10.89.5.17 netmask: 255.255.255.192 gateway: 10.89.5.1 on management0
```

```
Updating routing tables, please wait...
All configurations applied to the system. Took 3 Seconds.
Saving a copy of running network configuration to local disk.
For HTTP Proxy configuration, run 'configure network http-proxy'

Manage the device locally? (yes/no) [yes]: no
DHCP server is already disabled
DHCP Server Disabled
Configure firewall mode? (routed/transparent) [routed]:
Configuring firewall mode ...


Device is in OffBox mode - disabling/removing port 443 from iptables.
Update policy deployment information
    - add device configuration
    - add network discovery
    - add system policy

You can register the sensor to a Firepower Management Center and use the
Firepower Management Center to manage it. Note that registering the sensor
to a Firepower Management Center disables on-sensor Firepower Services
management capabilities.

When registering the sensor to a Firepower Management Center, a unique
alphanumeric registration key is always required.  In most cases, to register
a sensor to a Firepower Management Center, you must provide the hostname or
the IP address along with the registration key.
'configure manager add [hostname | ip address ] [registration key ]'

However, if the sensor and the Firepower Management Center are separated by a
NAT device, you must enter a unique NAT ID, along with the unique registration
key.
'configure manager add DONTRESOLVE [registration key ] [ NAT ID ]'

Later, using the web interface on the Firepower Management Center, you must
use the same registration key and, if necessary, the same NAT ID when you add
this sensor to the Firepower Management Center.
>
```

Step 5 Identify the Firewall Management Center that will manage this Firewall Threat Defense.

configure manager add {*hostname* | *IPv4_address* | *IPv6_address* | **DONTRESOLVE**} *reg_key* [*nat_id*] [*display_name*]

Note
If you are using Security Cloud Control for management, use the Security Cloud Control-generated **configure manager add** command for this step.

- {*hostname* | *IPv4_address* | *IPv6_address* | **DONTRESOLVE**}—Specifies either the FQDN or IP address of the Firewall Management Center. If the Firewall Management Center is not directly addressable, use **DONTRESOLVE** and also specify the *nat_id*. At least one of the devices, either the Firewall Management Center or the Firewall Threat Defense, must have a reachable IP address to establish the two-way, TLS-1.3-encrypted communication channel between the two devices. If you specify **DONTRESOLVE** in this command, then the Firewall Threat Defense must have a reachable IP address or hostname.
- *reg_key*—Specifies a one-time registration key of your choice that you will also specify on the Firewall Management Center when you register the Firewall Threat Defense. The registration key must be between 2 and 36 characters. Valid characters include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-).

- *nat_id*—Specifies a unique, one-time string of your choice that you will also specify on the Firewall Management Center when you register the Firewall Threat Defense. The NAT ID must be between 2 and 36 characters. Valid characters include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-). This ID cannot be used for any other devices registering to the Firewall Management Center. The NAT ID is used in combination with the IP address to verify that the connection is coming from the correct device; only after authentication of the IP address/NAT ID will the registration key be checked. We recommended that you always use the NAT ID even when it is optional, but it is required if:
  - You set the Firewall Management Center IP address to **DONTRESOLVE**.
  - When adding the device on the Firewall Management Center, you do not specify a reachable device IP address or hostname.
  - You use the data interface for management, even if you specify IP addresses on both sides.
  - The Firewall Management Center uses multiple management interfaces.
- *display_name*—Provide a display name for showing this manager with the **show managers** command. This option is useful if you are identifying Security Cloud Control as the primary manager and an on-prem Firewall Management Center for analytics only. If you don't specify this argument, the firewall auto-generates a display name using one of the following methods:
  - *hostname* | *IP_address* (if you don't use the **DONTRESOLVE** keyword)
  - **manager-***timestamp*

Example:

```
> configure manager add MC.example.com 123456
Manager successfully configured.
```

Example:

If the Firewall Management Center is behind a NAT device, enter a unique NAT ID along with the registration key, and specify DONTRESOLVE instead of the hostname, for example:

```
> configure manager add DONTRESOLVE regk3y78 natid90
Manager successfully configured.
```

Example:

If the Firewall Threat Defense is behind a NAT device, enter a unique NAT ID along with the Firewall Management Center IP address or hostname, for example:

```
> configure manager add 10.70.45.5 regk3y78 natid56
Manager successfully configured.
```

Step 6 If you are using Security Cloud Control as your primary manager and want to use an on-prem Firewall Management Center for analytics only, identify the on-prem Firewall Management Center.

configure manager add {*hostname* | *IPv4_address* | *IPv6_address* | **DONTRESOLVE**} *reg_key* [*nat_id*] [*display_name*]

Example:

The following example uses the generated command for Security Cloud Control with a Security Cloud Control-generated display name and then specifies an on-prem Firewall Management Center for analytics only with the "analytics-FMC" display name.

```
> configure manager add account1.app.us.cdo.cisco.com KPOOP0rgWzaHrnj1V5ha2q5Rf8pKFX9E
Lzm1HOynhVUWhXYWz2swmkj2ZWsN3Lb account1.app.us.cdo.cisco.com
Manager successfully configured.
> configure manager add 10.70.45.5 regk3y78 natid56 analytics-FMC
Manager successfully configured.
```

Step 7 (Optional) Configure a data interface for manager access.

configure network management-data-interface

You are then prompted to configure basic network settings for the data interface.

Note
You should use the console port when using this command. If you use SSH to the Management interface, you might get disconnected and have to reconnect to the console port. See below for more information about SSH usage.

See the following details for using this command. See also Using the Firewall Threat Defense Data Interface for Management, on page 4.

- The original Management interface cannot use DHCP if you want to use a data interface for management. If you did not set the IP address manually during initial setup, you can set it now using the **configure network** {**ipv4** | **ipv6**} **manual** command. Make sure this interface is on a different subnet from the manager access interface to prevent routing issues. If you did not already set the Management interface gateway to **data-interfaces**, this command will set it now.
- When you add the Firewall Threat Defense to the Firewall Management Center, the Firewall Management Center discovers and maintains the interface configuration, including the following settings: interface name and IP address, static route to the gateway, DNS servers, and DDNS server. For more information about the DNS server configuration, see below. In the Firewall Management Center, you can later make changes to the manager access interface configuration, but make sure you don't make changes that can prevent the Firewall Threat Defense or Firewall Management Center from re-establishing the management connection. If the management connection is disrupted, the Firewall Threat Defense includes the **configure policy rollback** command to restore the previous deployment.
- DDNS ensures the Firewall Management Center can reach the Firewall Threat Defense at its Fully-Qualified Domain Name (FQDN) if the IP address changes. If you configure a DDNS server update URL, the Firewall Threat Defense automatically adds certificates for all of the major CAs from the Cisco Trusted Root CA bundle so that the Firewall Threat Defense can validate the DDNS server certificate for the HTTPS connection. The Firewall Threat Defense supports any DDNS server that uses the DynDNS Remote API specification (https://help.dyn.com/remote-access-api/).
- This command sets the *data* interface DNS server. The Management DNS server that you set with the setup script (or using the **configure network dns servers** command) is used for management traffic. The data DNS server is used for DDNS (if configured) or for security policies applied to this interface.

  On the Firewall Management Center, the data interface DNS servers are configured in the Platform Settings policy that you assign to this Firewall Threat Defense. When you add the Firewall Threat Defense to the Firewall Management Center, the local setting is maintained, and the DNS servers are *not* added to a Platform Settings policy. However, if you later assign a Platform Settings policy to the Firewall Threat Defense that includes a DNS configuration, then that configuration will overwrite the local setting. We suggest that you actively configure the DNS Platform Settings to match this setting to bring the Firewall Management Center and the Firewall Threat Defense into sync.

  Also, local DNS servers are only retained by the Firewall Management Center if the DNS servers were discovered at initial registration. For example, if you registered the device using the Management interface,

but then later configure a data interface using the **configure network management-data-interface** command, then you must manually configure all of these settings in the Firewall Management Center, including the DNS servers, to match the FTD configuration.

- You can change the management interface after you register the Firewall Threat Defense to the Firewall Management Center, to either the Management interface or another data interface.
- The FQDN that you set in the setup wizard will be used for this interface.
- You can clear the entire device configuration as part of the command; you might use this option in a recovery scenario, but we do not suggest you use it for initial setup or normal operation.
- To disable data management, enter the **configure network management-data-interface disable** command.

Example:

```
> configure network management-data-interface
Data interface to use for management: ethernet1/1
Specify a name for the interface [outside]:
IP address (manual / dhcp) [dhcp]:
DDNS server update URL [none]:
https://dwinchester:pa$$w0rd17@domains.example.com/nic/update?hostname=<h>&myip=<a>
Do you wish to clear all the device configuration before applying ? (y/n) [n]:

Configuration done with option to allow manager access from any network,
if you wish to change the manager access network
use the 'client' option in the command 'configure network management-data-interface'.

Setting IPv4 network configuration.
Network settings changed.

>
```

Example:

```
> configure network management-data-interface
Data interface to use for management: ethernet1/1
Specify a name for the interface [outside]: internet
IP address (manual / dhcp) [dhcp]: manual
IPv4/IPv6 address: 10.10.6.7
Netmask/IPv6 Prefix: 255.255.255.0
Default Gateway: 10.10.6.1
Comma-separated list of DNS servers [none]: 208.67.222.222,208.67.220.220
DDNS server update URL [none]:
Do you wish to clear all the device configuration before applying ? (y/n) [n]:

Configuration done with option to allow manager access from any network,
if you wish to change the manager access network
use the 'client' option in the command 'configure network management-data-interface'.

Setting IPv4 network configuration.
Network settings changed.

>
```

Step 8 (Optional) Limit data interface access to a manager on a specific network.

configure network management-data-interface client *ip_address netmask*

By default, all networks are allowed.

What to do next

Register your device to a Firewall Management Center.

Configure an Event Interface

You always need a management interface for management traffic. If your device has a second management interface, for example, the Firepower 4100/9300 and Secure Firewall 4200, you can enable it for event-only traffic.

Before you begin

To use a separate event interface, you also need to enable an event interface on the Firewall Management Center. See the Cisco Secure Firewall Management Center Administration Guide.

Procedure

Step 1 Enable the second management interface as an event-only interface.

configure network management-interface enable management1

configure network management-interface disable-management-channel management1

You can optionally disable events for the main management interface using the **configure network management-interface disable-events-channel** command. In either case, the device will try to send events on the event-only interface, and if that interface is down, it will send events on the management interface even if you disable the event channel.

You cannot disable both event and management channels on an interface.

Example:

```
> configure network management-interface enable management1
Configuration updated successfully

> configure network management-interface disable-management-channel management1
Configuration updated successfully

>
```

Step 2 Configure the IP address of the event interface.

The event interface can be on a separate network from the management interface, or on the same network.

a) Configure the IPv4 address:

configure network ipv4 manual *ip_address netmask gateway_ip* **management1**

Note that the *gateway_ip* in this command is used to create the default route for the device, so you should enter the value you already set for the management0 interface. It does not create a separate static route

for the eventing interface. If you are using an event-only interface on a different network from the management interface, we recommend that you create a static route separately for the event-only interface.

Example:

```
> configure network ipv4 manual 10.10.10.45 255.255.255.0 10.10.10.1 management1
Setting IPv4 network configuration.
Network settings changed.

>
```

b) Configure the IPv6 address:

- Stateless autoconfiguration:

  configure network ipv6 router management1

  Example:

  ```
  > configure network ipv6 router management1
  Setting IPv6 network configuration.
  Network settings changed.

  >
  ```

- Manual configuration:

  configure network ipv6 manual *ip6_address ip6_prefix_length* **management1**

  Example:

  ```
  > configure network ipv6 manual 2001:0DB8:BA98::3210 64 management1
  Setting IPv6 network configuration.
  Network settings changed.

  >
  ```

Step 3 Add a static route for the event-only interface if the Firewall Management Center is on a remote network; otherwise, all traffic will match the default route through the management interface.

configure network static-routes {**ipv4** | **ipv6**}**add management1** *destination_ip netmask_or_prefix gateway_ip*

For the *default* route, do not use this command; you can only change the default route gateway IP address when you use the **configure network ipv4** or **ipv6** commands (see, Step 2, on page 26).

Example:

```
> configure network static-routes ipv4 add management1 192.168.6.0 255.255.255.0 10.10.10.1
Configuration updated successfully

> configure network static-routes ipv6 add management1 2001:0DB8:AA89::5110 64
2001:0DB8:BA98::3211
Configuration updated successfully

>
```

To display static routes, enter **show network-static-routes** (the default route is not shown):

```
> show network-static-routes
---------------[ IPv4 Static Routes ]---------------
Interface                 : management1
Destination               : 192.168.6.0
Gateway                   : 10.10.10.1
Netmask                   : 255.255.255.0
[…]
```

Manage Devices

Register and unregister devices to the Firewall Management Center.

About the Device Management Page

The **Devices** > **Device Management** page provides you with range of information and options.

Figure 9: Device Management Page

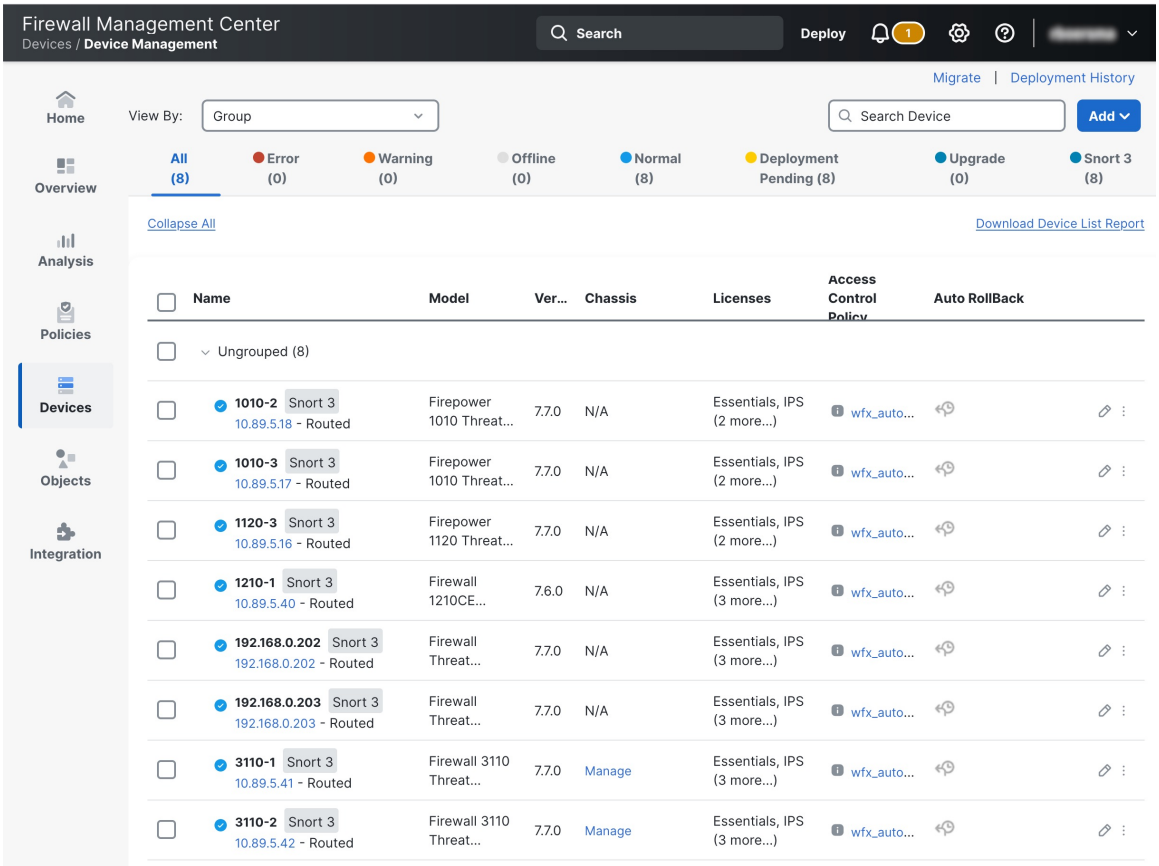

- **View By**—View devices based on group, licenses, model, version, or access control policy.
- **Device State**—View devices based on state (**Error**, **Warning**, etc.). You can click on a state icon to view the devices belonging to it. The number of devices belonging to the states are provided within brackets.
- **Search Device**—Search for a device by device name, host name, or IP address.
- **Add**—Add devices and other manageable components.

Figure 10: Add Menu

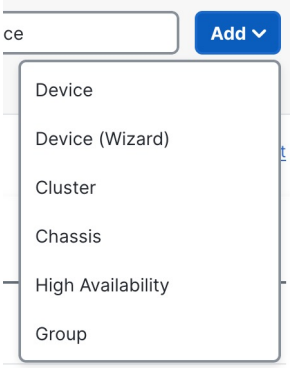

- Columns—Click the column head to sort by that column.
  - **Name**
  - **Model**
  - **Version**
  - **Chassis**—For supported models, click **Manage** to bring up the integrated Chassis Manager. For the Firepower 4100/9300, the link cross-launches the Firewall Chassis Manager.
  - **Licenses**
  - **Access Control Policy**—Click on the link in the Access Control Policy column to view the policy that is deployed to the device.
  - **Auto-Rollback**—Shows whether auto-rollback of the configuration is enabled ( ) or disabled ( ) if the deployment causes the management connection to go down. See Edit Deployment Settings, on page 201.
- **Edit**—For each device, use the **Edit** ( ) icon to edit the device settings.

  You can also just click on the device name or IP address.
- **More**—For each device, click the **More** ( ) icon to execute other actions:

Figure 11: More Menu

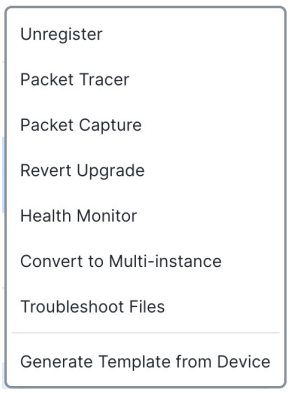

- **Unregister**—To unregister the device.
- **Packet Tracer**—To navigate to the packet tracer page for examining policy configuration on the device by injecting a model packet into the system.
- **Packet Capture**—To navigate to the packet capture page, where, you can view the verdicts and actions the system takes while processing a packet.
- **Revert Upgrade**—To revert the upgrade and configuration changes that were made after the last upgrade. This action results in restoring the device to the version that was before the upgrade.
- **Health Monitor**—To navigate to the device's health monitoring page.
- **Convert to Multi-instance**—For supported models, convert the chassis to multi-instance mode.
- **Troubleshoot Files**—Generate troubleshooting files, where you can choose the type of data to be included in the report.
- **Generate Template from Device**—Generate a new device template from a registered device. The new template has the same configuration as the device from which it is generated. You can generate a new device template from standalone and HA devices. However, if you generate a template from HA devices, the new template will not contain the failover configurations.

Add a Device Group

The Firewall Management Center allows you to group devices so you can easily deploy policies and install updates on multiple devices. You can expand and collapse the list of devices in the group.

If you add the primary device in a high-availability pair to a group, both devices are added to the group. If you break the high-availability pair, both devices remain in that group.

Groups are not supported in a multidomain environment.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 From the **Add** drop-down menu, choose **Add Group**.

To edit an existing group, click **Edit** (✎) for the group you want to edit.

Step 3 Enter a **Name**.

Step 4 Under **Available Devices**, choose one or more devices to add to the device group. Use Ctrl or Shift while clicking to choose multiple devices.

Step 5 Click **Add** to include the devices you chose in the device group.

Step 6 Optionally, to remove a device from the device group, click **Delete** (🗑) next to the device you want to remove.

Step 7 Click **OK** to add the device group.

Register With the Management Center

The Firewall Management Center offers multiple methods to register your devices.

Registration Key Method

Add a device using a registration key that you specify in both the Firewall Management Center and the device initial configuration.

Add a Device Using a Registration Key—Basic Configuration

Use this procedure to add a device to the Firewall Management Center using a registration key and a basic configuration; to use a device template, see Add a Device Using a Registration Key—Device Template, on page 36. If you plan to link devices for high availability, you must still use this procedure; see Add a High Availability Pair, on page 441. For clustering, see the clustering chapter for your model.

You can also use this procedure to add a device that is managed by a Cloud-Delivered Firewall Management Center, and you want to use the on-prem Firewall Management Center for event logging and analytics purposes only.

If you use Firewall Management Center high availability, add devices *only* to the active Firewall Management Center. Devices registered to the active Firewall Management Center are automatically registered to the standby.

Before you begin

- Set up the device to be managed by the Firewall Management Center. See:
  - Complete the Firewall Threat Defense Initial Configuration for Manual Registration, on page 13
  - The getting started guide for your model
- The Firewall Management Center must be registered to the Smart Software Manager. A valid evaluation license is sufficient, but if it expires, you will not be able to add new devices until you successfully register.
- If you registered a device using IPv4 and want to convert it to IPv6, you must delete and reregister the device.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 From the **Add** drop-down menu, choose **Device (Wizard)**.

Step 3 Click **Registration Key**, and then click **Next**.

Figure 12: Device Registration Method

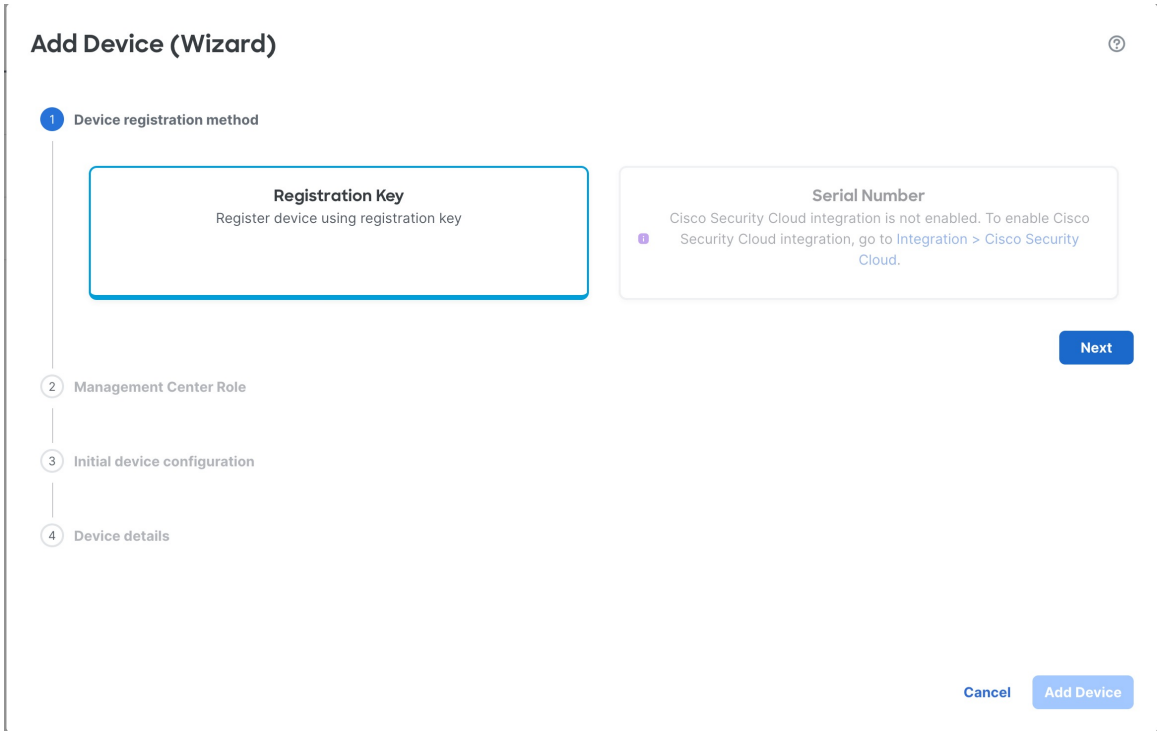

Step 4 In a multi-domain environment, choose the **Domain** from the drop-down list and click **Next**.

Figure 13: Domain

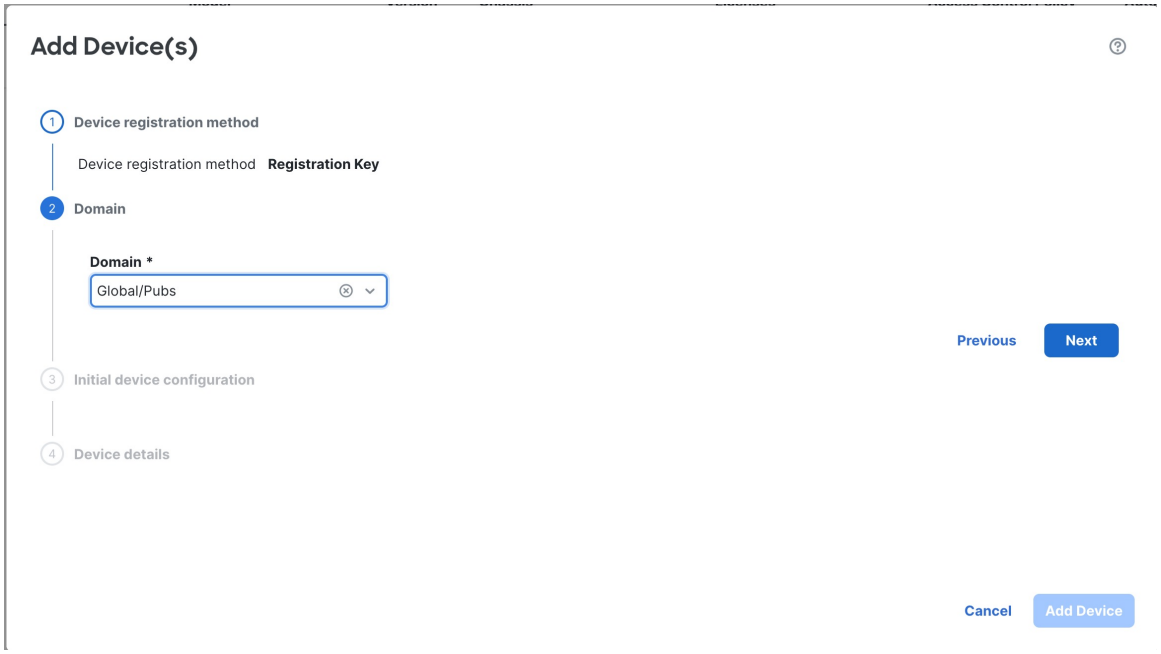

Step 5 Click **Primary manager** for normal management or **Analytics-only manager** for a device that is managed by a Cloud-Delivered Firewall Management Center. Analytics-only mode does not support a multi-domain environment, so this step doesn't appear in that case.

Figure 14: Management Center Role

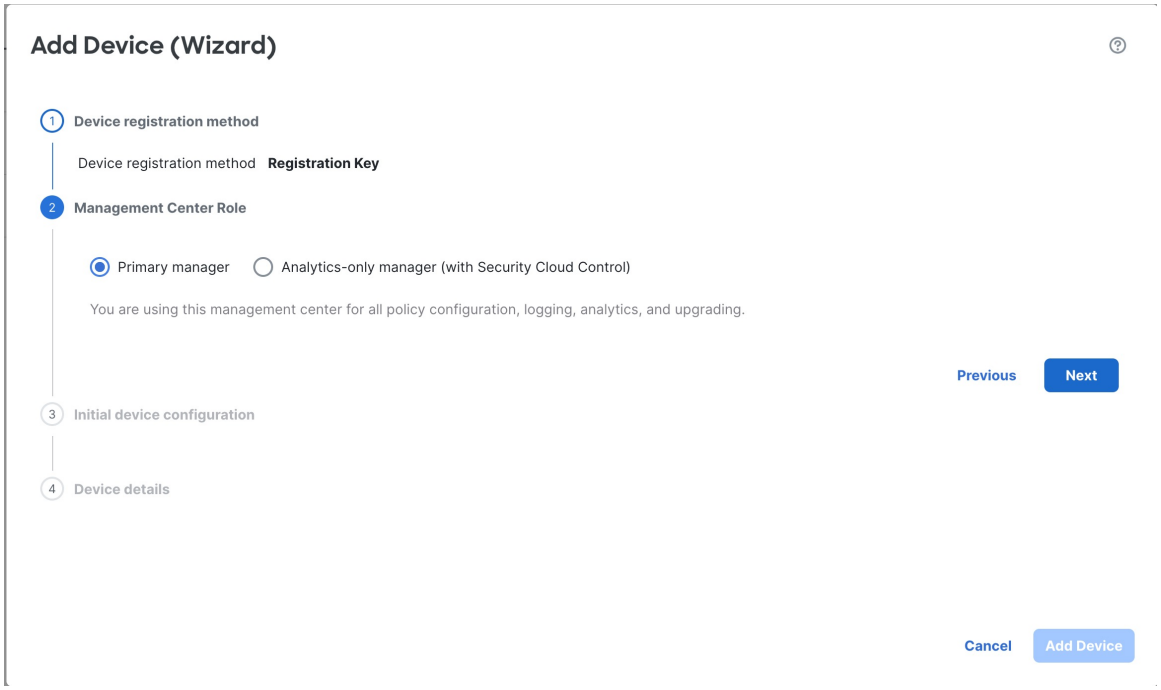

Step 6 For the **Initial Device Configuration**, click **Basic**.

Figure 15: Initial Device Configuration

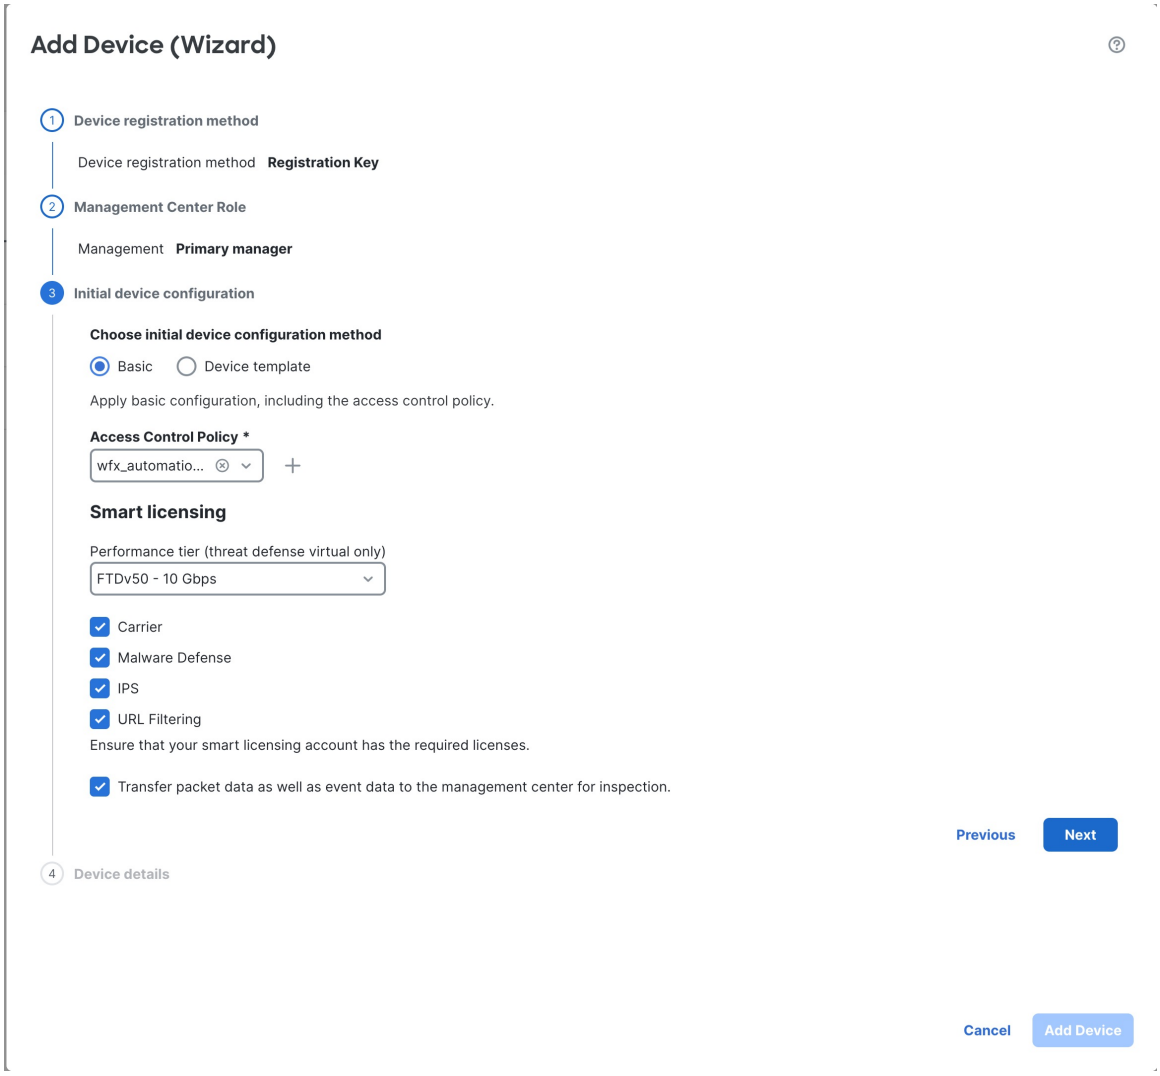

For analytics-only mode, the system hides **Initial Device Configuration** because these settings are managed by Security Cloud Control.

a) Choose an initial **Access Control Policy** to deploy to the device at registration, or create a new policy.
b) Choose **Smart Licensing** licenses to apply to the device.

   You can also apply licenses after you add the device, from the **System** > **Licenses** > **Smart Licenses** page, including the Secure Client remote access VPN license.

   For the Firewall Threat Defense Virtual only, you must also select the **Performance Tier**. It's important to choose the tier that matches the license you have in your account. Until you choose a tier, your device defaults to FTDv50.

c) Click **Next**.

Step 7 Specify the **Device details**.

Figure 16: Device Details

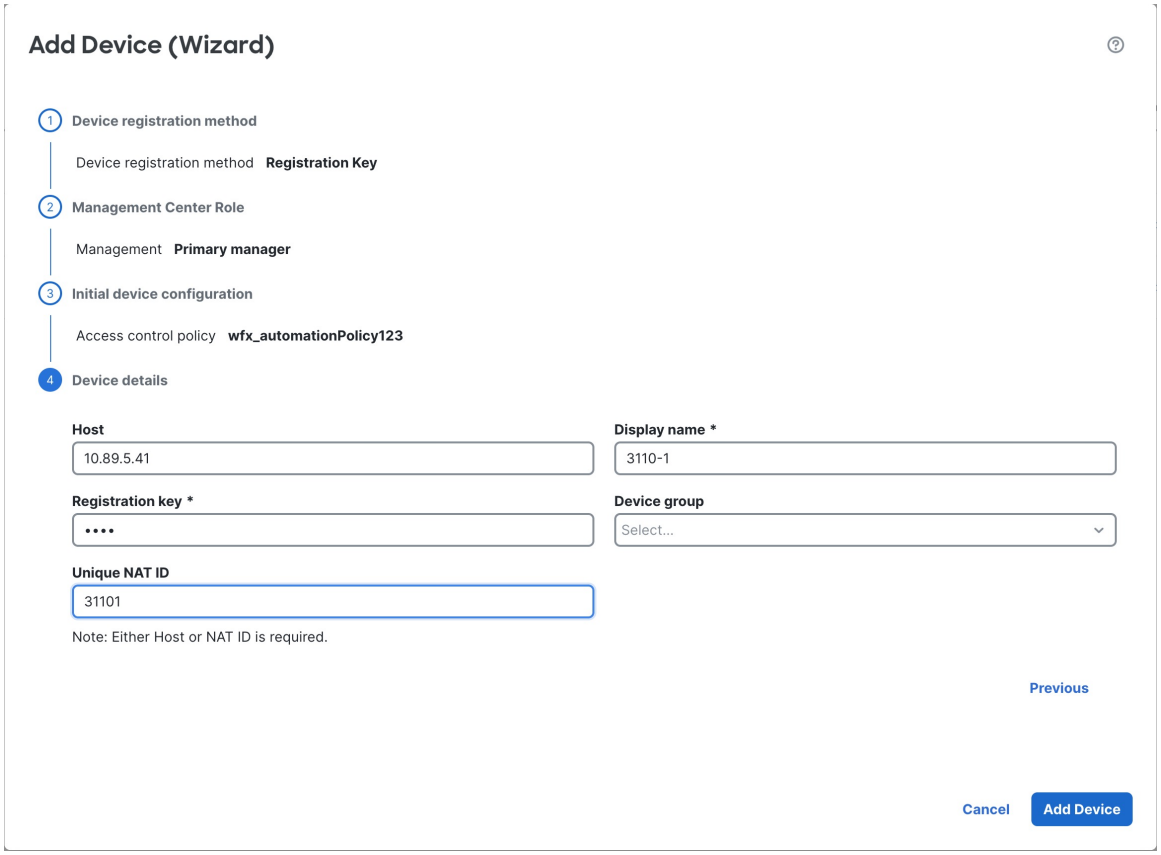

a) For the **Host**, enter the IP address or the hostname of the device you want to add. Leave this field blank if you don't know the device IP address (for example, it's behind NAT).

   If you leave this field blank, the initial configuration on the device needs to include a reachable Firewall Management Center IP address or hostname plus the NAT ID. For more information, see NAT Environments, on page 7.

b) For the **Display name**, enter a name for the device as you want it to display in the Firewall Management Center. You cannot change this name later.

c) For the **Registration Key**, enter the same registration key from your initial configuration. The registration key is a one-time-use shared secret. The key can be up to 37-characters in length and include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-). The registration key does not need to be unique per device.

d) (Optional) Add the device to a **Device group**

e) For the **Unique NAT ID**, enter the same ID from your initial configuration.

   The **Unique NAT ID** specifies a unique, one-time string of your choice that you will also specify on the device during initial configuration. It is required when one side does not specify a reachable IP address or hostname, for example if you left the **Host** field blank. Although technically optional, we recommend always specifying the NAT ID even when you know the IP addresses of both sides because it is required in certain situations. The ID can be up to 37-characters in length and include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-). This ID cannot be used for any other devices registering to the Firewall Management Center.

f) Check **Transfer Packets** so that for each intrusion event, the device transfers the packet to the Firewall Management Center for inspection.

For each intrusion event, the device sends event information and the packet that triggered the event to the Firewall Management Center for inspection. If you disable it, only event information will be sent to the Firewall Management Center; the packet will not be sent.

Step 8 Click **Add Device**.

It may take up to two minutes for the Firewall Management Center to verify the device's heartbeat and establish communication. If the registration succeeds, the device is added to the list. If it fails, you will see an error message. If the device fails to register, check the following items:

- Ping—Access the device CLI, and ping the Firewall Management Center IP address using the following command:

  ping system *ip_address*

  If the ping is not successful, check your network settings using the **show network** command. If you need to change the device IP address, use the **configure network** {**ipv4** | **ipv6**} **manual** command.

- Registration key, NAT ID, and Firewall Management Center IP address—Make sure you are using the same registration key, and if used, NAT ID, on both devices. You can set the registration key and NAT ID on the device using the **configure manager add** command.

For more troubleshooting information, see https://cisco.com/go/fmc-reg-error.

Add a Device Using a Registration Key—Device Template

You can use a template to add a device, register the device with the Firewall Management Center and bring up the device with the given template configurations.

Before you begin

Create a device template according to Device Management Using Device Templates, on page 89. You must specify any required variables and network object overrides for each device and ensure that model mapping is done for the target device model.

We recommend that you create a checklist to ensure that all configurations in the template have been entered correctly before applying the template on the device.

A sample checklist is given below.

- Check version, model, operation modes.
- Check list of variables and overrides.
- Check sanity of variable and override values.
- Check if the required Model Mappings exist.
- Check if parallel device template operations are in progress.

Note If you are adding a device that will be managed by a data interface, ensure that you configure the template to be compatible with the connectivity parameters of the device. For more information, see Configure a Template for Threat Defense Devices Managed Through the Data Interface.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 From the **Add** drop-down menu, choose **Device (Wizard)**.

Step 3 Click **Registration Key**, and then click **Next**.

Figure 17: Device Registration Method

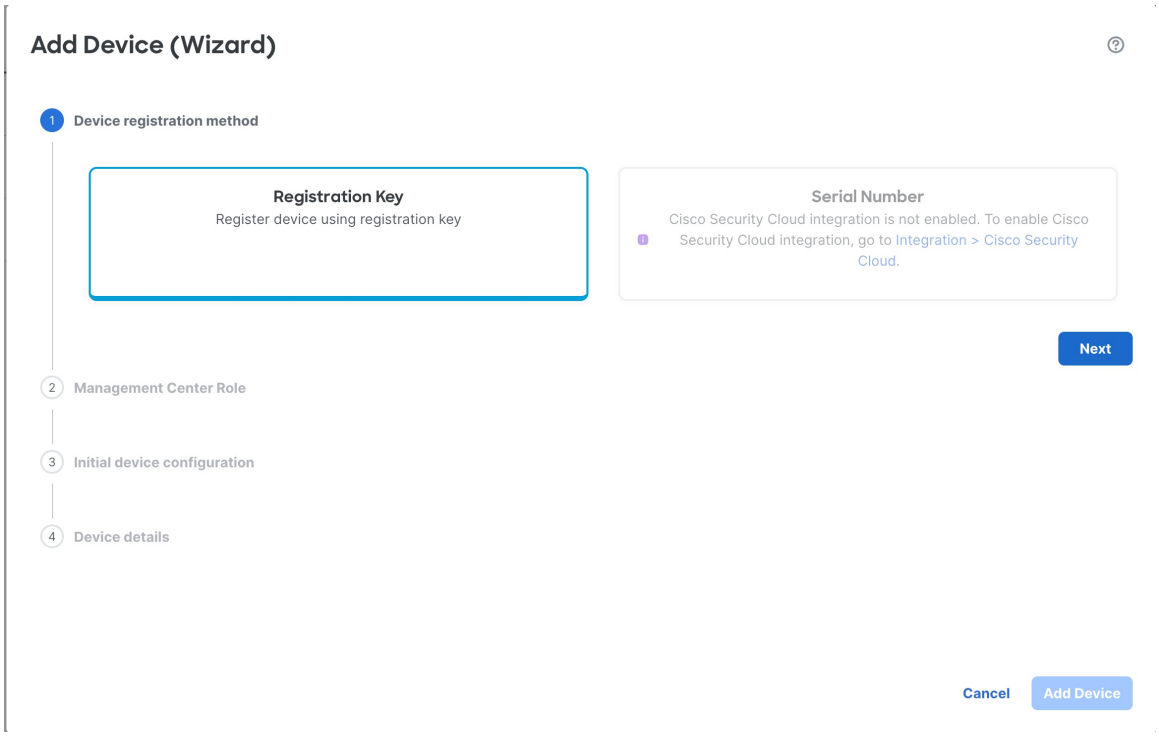

Step 4 In a multi-domain environment, choose the **Domain** from the drop-down list and click **Next**.

Figure 18: Domain

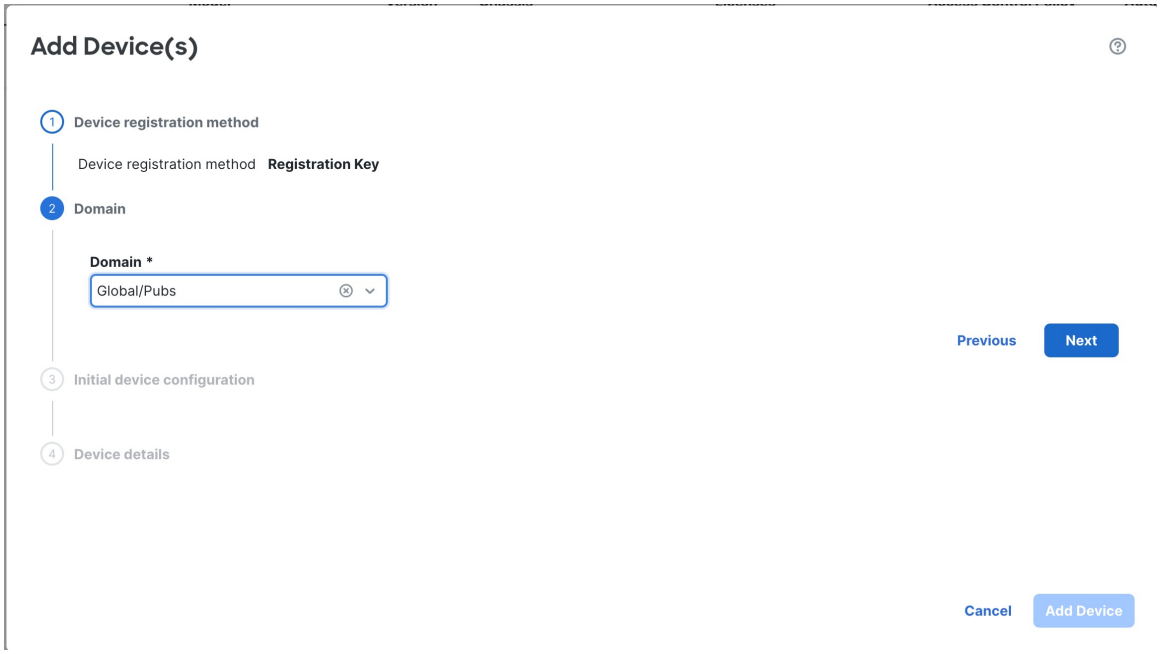

Step 5 In **Initial device configuration**, configure the following settings.

Figure 19: Initial Device Configuration

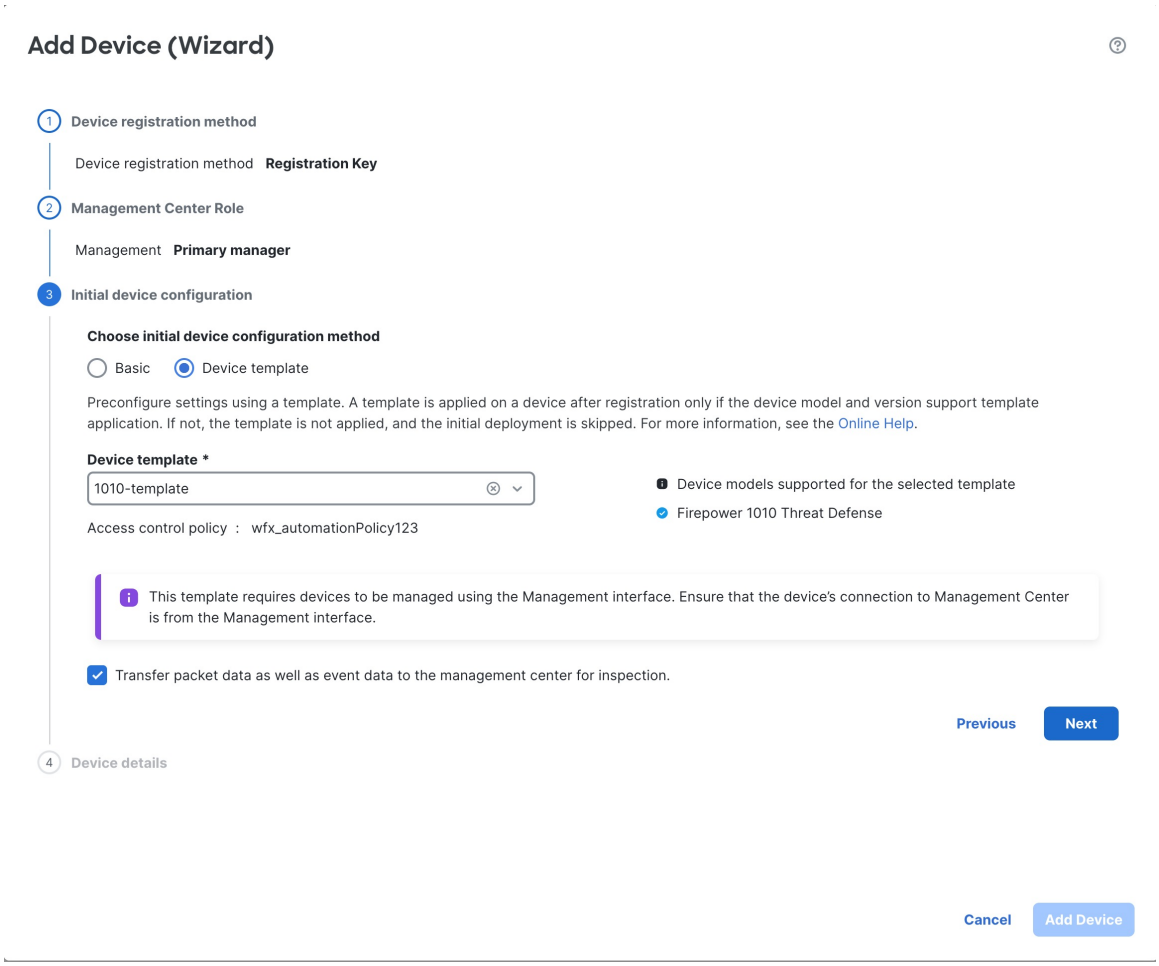

a) Click **Device template**.
b) Choose a template from the **Device template** drop-down list.
c) (Optional) Click **Transfer packet data** so that for each intrusion event, the device transfers the packet to the Firewall Management Center for inspection.

   For each intrusion event, the device sends event information and the packet that triggered the event to the Firewall Management Center for inspection. If you disable it, only event information will be sent to the Firewall Management Center; the packet will not be sent.

d) Click **Next**.

Step 6 Specify the **Device details**.

Figure 20: Device Details

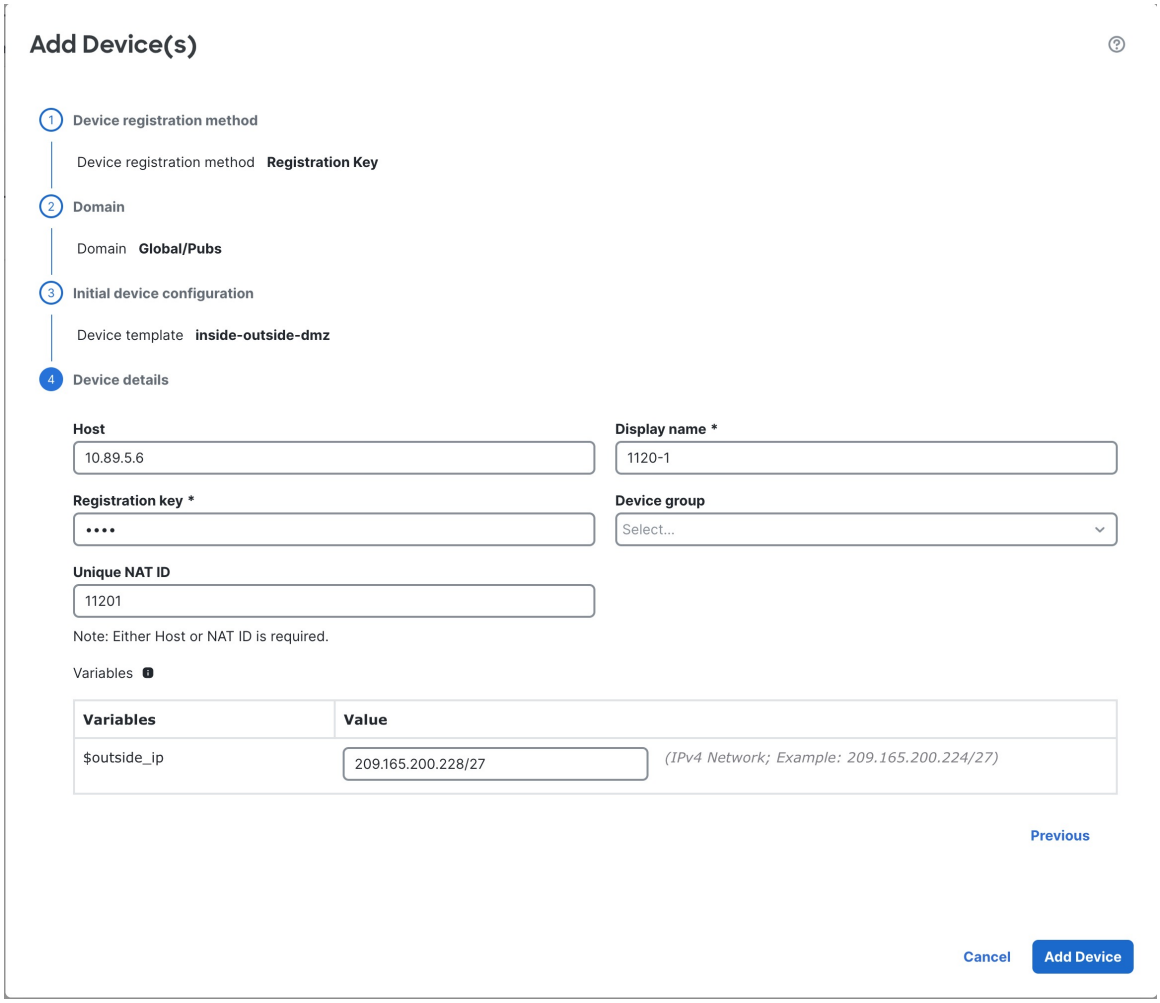

a) For the **Host**, enter the IP address or the hostname of the device you want to add. Leave this field blank if you don't know the device IP address (for example, it's behind NAT).

   If you leave this field blank, the initial configuration on the device needs to include a reachable Firewall Management Center IP address or hostname plus the NAT ID. For more information, see NAT Environments, on page 7.

b) For the **Display name**, enter a name for the device as you want it to display in the Firewall Management Center. You cannot change this name later.

c) For the **Registration Key**, enter the same registration key in your initial configuration. The registration key is a one-time-use shared secret. The key can be up to 37-characters in length and include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-). The registration key does not need to be unique per device.

d) (Optional) Add the device to a **Device group**

e) If you specified a NAT ID during initial configuration, enter the same NAT ID for the **Unique NAT ID**.

   The **Unique NAT ID** specifies a unique, one-time string of your choice that you will also specify on the device during initial configuration. It is required when one side does not specify a reachable IP address or hostname, for example if you left the **Host** field blank. Although technically optional, we recommend always specifying the NAT ID even when you know the IP addresses of both sides because it is required

in certain situations. The ID can be up to 37-characters in length and include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-). This ID cannot be used for any other devices registering to the Firewall Management Center.

f) Check **Transfer Packets** so that for each intrusion event, the device transfers the packet to the Firewall Management Center for inspection.

For each intrusion event, the device sends event information and the packet that triggered the event to the Firewall Management Center for inspection. If you disable it, only event information will be sent to the Firewall Management Center; the packet will not be sent.

g) Enter values for the **Variables** and **Network object overrides**.

Step 7 Click **Add Device** to initiate device registration. The template configurations are applied after the device is successfully registered with the Firewall Management Center.

Serial Number Method (Zero-Touch Provisioning)

Zero-Touch Provisioning lets you register devices to the Firewall Management Center by serial number without having to perform any initial setup on the device.

Add a Device Using the Serial Number (Zero-Touch Provisioning)—Basic Configuration

Zero-Touch Provisioning lets you register devices to the Firewall Management Center by serial number without having to perform any initial setup on the device. The Firewall Management Center integrates with the Cisco Security Cloud and Security Cloud Control for this functionality.

Use this procedure to add a single device to the Firewall Management Center using a basic configuration. To add one or more devices using a template, see Add Devices Using Serial Numbers (Zero-Touch Provisioning)—Device Template.

Default Configuration

When you use zero-touch provisioning, the following interfaces are preconfigured. Note that other settings, such as the DHCP server on inside, access control policy, or security zones, are not configured.

- Ethernet 1/1—"outside", IP address from DHCP, IPv6 autoconfiguration
- Ethernet 1/2 (or for the 1010/ 1210// 1220, the VLAN1 interface)— "inside", 192.168.95.1/24
- Default route—Obtained through DHCP on the outside interface

Requirements

Zero-Touch Provisioning is not supported with clustering or multi-instance mode.

When you use the outside interface for manager access, it uses DHCP by default. Before you can enable high availability, you need to change the IP address to a static address. See Change the Device IP Address, on page 158. Alternatively, you can use the Management interface instead; DHCP is supported on Management with high availability.

Zero-Touch Provisioning is only supported on the following models using 7.2 and 7.4 or later; prior to 7.2.4, the Firewall Management Center must be publicly reachable.

- Firepower 1010
- Firepower 1100

- Secure Firewall 1200
- Firepower 2100 (on supported device versions)
- Secure Firewall 3100

Before you begin

- Make sure the device is unconfigured or a fresh install. Zero-Touch Provisioning is meant for new devices only. Pre-configuration can disable zero-touch provisioning, depending on how you configure the device. If you disable zero-touch provisioning, see Resolve Serial Number (Zero-Touch Provisioning) Registration Issues, on page 57.
- Cable the outside interface or Management interface so it can reach the internet. If you use the outside interface for zero-touch provisioning, do not also cable the Management interface; if the Management interface gets an IP address from DHCP, the routing will be incorrect for the outside interface.
- If the device does not have a public IP address or FQDN, or you use the Management interface, set a public IP address/FQDN for the Firewall Management Center (for example, if it is behind NAT), so the device can initiate the management connection. See **System** > **Configuration** > **Manager Remote Access**.
- The Firewall Management Center must be registered to the Smart Software Manager. A valid evaluation license is sufficient, but if it expires, you will not be able to add new devices until you successfully register.
- If you registered a device using IPv4 and want to convert it to IPv6, you must unregister and reregister the device.

Procedure

Step 1 The first time you add a device using a serial number, integrate the Firewall Management Center with Cisco Security Cloud.

Note
For a Firewall Management Center high-availability pair, you also need to integrate the secondary Firewall Management Center with Cisco Security Cloud.

a) Choose **Integration** > **Cisco Security Cloud**.
b) Click **Enable Cisco Security Cloud** to open a separate browser tab to log you into your Cisco Security Cloud account and confirm the displayed code.

Make sure this page is not blocked by a pop-up blocker. If you do not already have a Cisco Security Cloud and Security Cloud Control account, you can add one during this procedure.

For detailed information about this integration, see the "System Configuration" chapter in the Cisco Secure Firewall Management Center Administration Guide.

Security Cloud Control onboards the on-prem Firewall Management Center after you integrate the Firewall Management Center with Cisco Security Cloud. Security Cloud Control needs the Firewall Management Center in its inventory for zero-touch provisioning to operate. However, you do not need to use Security Cloud Control directly. If you do use Security Cloud Control, its Firewall Management Center support is

limited to device onboarding, viewing its managed devices, viewing objects associated with the Firewall Management Center, and cross-launching the Firewall Management Center.

c) Make sure **Enable Zero-Touch Provisioning** is checked.
d) Click **Save**.

Step 2 Obtain your device's serial number.

- If you have the shipping box, you can see the serial number on the label.
- The serial number is on a label on the bottom, or the back, or a pull-out tab at the front of the device.
- If you have console access, in FXOS, enter **show chassis detail**. Note that the correct serial number is called **Serial (SN)**. Do not use the **PCB Serial Number**. At the Firewall Threat Defense CLI, enter **show inventory** (not **show serial-number**, which shows the PCB serial number). Be careful not to disable zero-touch provisioning by entering certain settings at the Firewall Threat Defense startup script. If you disable zero-touch provisioning, see Resolve Serial Number (Zero-Touch Provisioning) Registration Issues, on page 57.

Step 3 Choose **Devices** > **Device Management**.

Step 4 From the **Add** drop-down menu, choose **Device (Wizard)**.

Step 5 Click **Use Serial Number**, and then click **Next**.

Figure 21: Device Registration Method

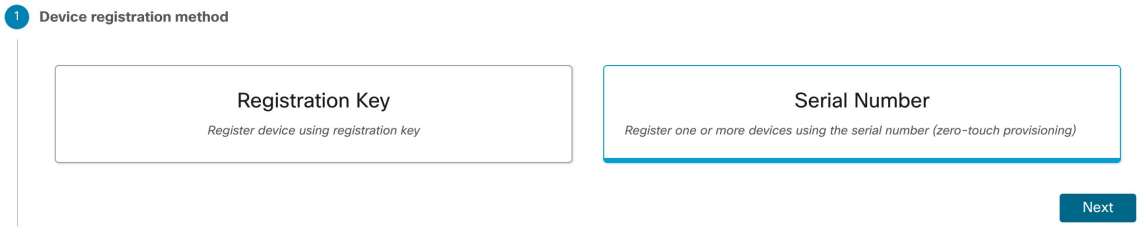

Step 6 In a multi-domain environment, choose the **Domain** from the drop-down list and click **Next**.

Figure 22: Domain

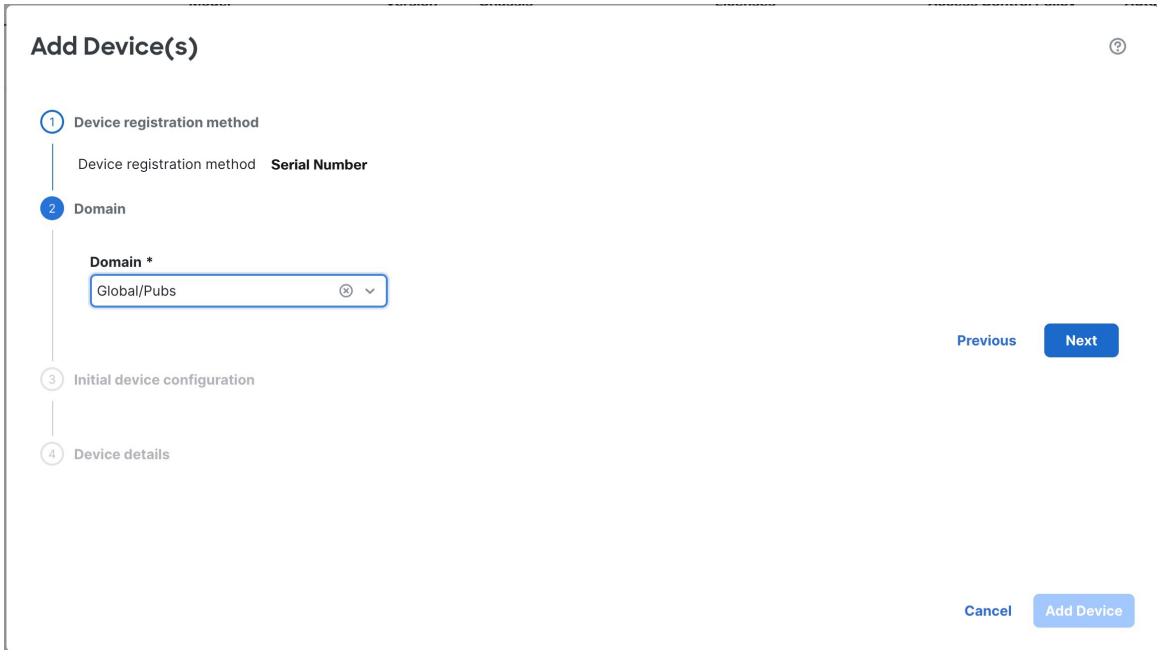

Step 7 For the **Initial device configuration**, click the **Basic** radio button.

Figure 23: Initial Device Configuration Method

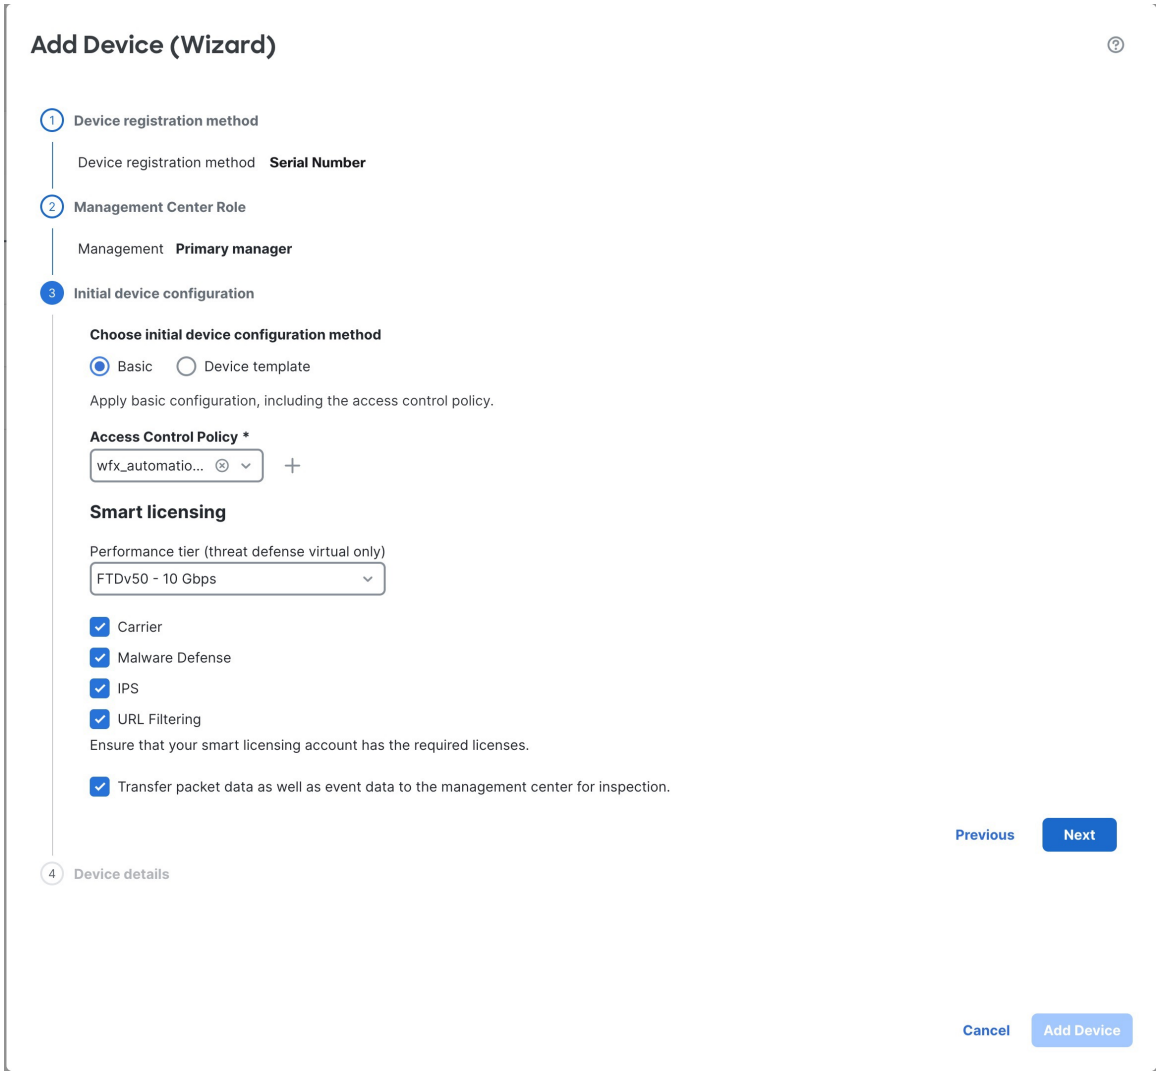

a) Choose an initial **Access Control Policy** to deploy to the device upon registration, or create a new policy.

If the device is incompatible with the policy you choose, deploying will fail. This incompatibility could occur for multiple reasons, including licensing mismatches, model restrictions, passive vs inline issues, and other misconfigurations. After you resolve the issue that caused the failure, manually deploy configurations to the device.

b) Choose **Smart licensing** licenses to apply to the device.

You can also apply licenses after you add the device, from the **System** > **Licenses** > **Smart Licenses** page.

c) Click **Next**.

Step 8 Configure the **Device details**.

Figure 24: Device details

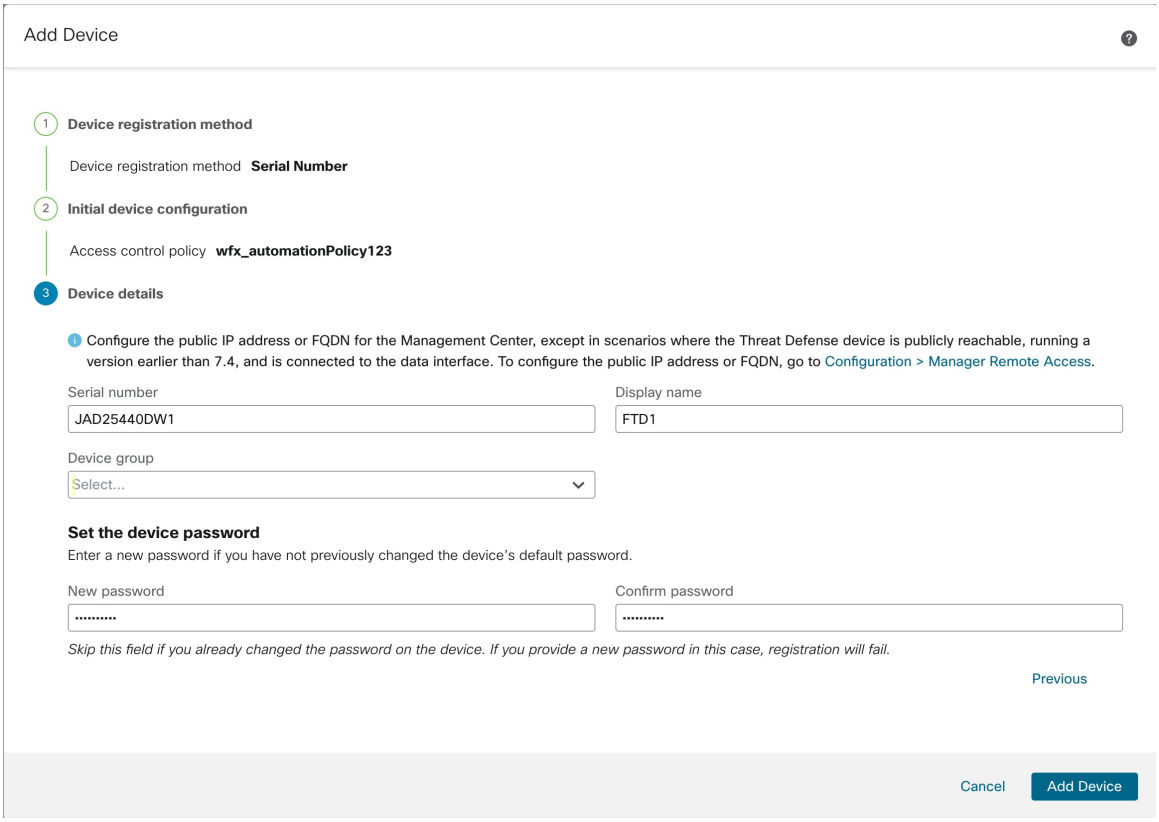

a) Enter the **Serial number**.
b) Enter the **Display name** as you want it to display in the Firewall Management Center
c) (Optional) Choose the **Device Group**.
d) **Set the device password**.

   If this device is unconfigured or a fresh install, then you need to set a new password. If you already logged in and changed the password, then leave this field blank. Otherwise, registration will fail.

Step 9 Click **Add Device**.

It may take up to two minutes for the Firewall Management Center to verify the device's heartbeat and establish communication. If the registration succeeds, the device is added to the list.

When using zero-touch provisioning on the outside interface, Security Cloud Control acts as a DDNS provider and does the following:

- Enables DDNS on outside using the **FMC Only** method. This method is only supported for zero-touch provisioning devices.
- Maps the outside IP address with the following hostname: *serial-number*.**local**.
- Provides the IP address/hostname mapping to the Firewall Management Center so it can resolve the hostname to the correct IP address.
- Informs the Firewall Management Center if the IP address ever changes, for example, if the DHCP lease renews.

If you use zero-touch provisioning on the Management interface, DDNS is not supported. The Firewall Management Center must be publicly reachable so the device can initiate the management connection.

You can continue to use Security Cloud Control as the DDNS provider, or you can later change the DDNS configuration in the Firewall Management Center to a different method. See Configure Dynamic DNS, on page 920 for more information.

If the device fails to register, see Resolve Serial Number (Zero-Touch Provisioning) Registration Issues, on page 57.

Add Devices Using Serial Numbers (Zero-Touch Provisioning)—Device Template

Zero-Touch Provisioning lets you register devices to the Firewall Management Center by serial number without having to perform any initial setup on the device. The Firewall Management Center integrates with the Cisco Security Cloud and Security Cloud Control for this functionality.

You can use a template to add a device, register the device with the Firewall Management Center and bring up the device with template configurations.

Use this procedure to add devices to the Firewall Management Center using serial numbers and a device template. To add a device without using a template, see Add a Device Using the Serial Number (Zero-Touch Provisioning)—Basic Configuration, on page 41.

Requirements

Zero-Touch Provisioning is not supported with clustering or multi-instance mode.

High availability is only supported when you use the Management interface because zero-touch provisioning uses DHCP, which is not supported for data interfaces and high availability.

Zero-Touch Provisioning with templates is supported on the following models using 7.4 or later:

- Firepower 1010
- Firepower 1100
- Secure Firewall 1200
- Firepower 2100 (on supported device versions)
- Secure Firewall 3100

Before you begin

- Make sure the device is unconfigured or a fresh install. Zero-Touch Provisioning is meant for new devices only. Pre-configuration can disable zero-touch provisioning, depending on how you configure the device. If you disable zero-touch provisioning, see Resolve Serial Number (Zero-Touch Provisioning) Registration Issues, on page 57.
- Cable the outside interface or Management interface so it can reach the internet. If you use the outside interface for zero-touch provisioning, do not also cable the Management interface; if the Management interface gets an IP address from DHCP, the routing will be incorrect for the outside interface.
- If the device does not have a public IP address or FQDN, or you use the Management interface, set a public IP address/FQDN for the Firewall Management Center (for example, if it is behind NAT), so the

device can initiate the management connection. See **System** > **Configuration** > **Manager Remote Access**.

- The Firewall Management Center must be registered to the Smart Software Manager. A valid evaluation license is sufficient, but if it expires, you will not be able to add new devices until you successfully register.
- If you registered a device using IPv4 and want to convert it to IPv6, you must unregister and reregister the device.
- Create a device template according to Device Management Using Device Templates, on page 89. You must specify any required variables and network-object overrides for each device and ensure that model mapping is done for the target device model.

  We recommend that you create a checklist to ensure that all configurations in the template have been entered correctly before applying the template on the device.

  A sample checklist is given below.

  - Check version, model, operation modes.
  - Check list of variables and overrides.
  - Check sanity of variable and override values.
  - Check if the required Model Mappings exist.
  - Check if parallel device template operations are in progress.

Note If you are adding a device that will be managed by a data interface, ensure that you configure the template to be compatible with the connectivity parameters of the device. For more information, see Configure a Template for Threat Defense Devices Managed Through the Data Interface.

Procedure

Step 1 The first time you add a device using a serial number, integrate the Firewall Management Center with Cisco Security Cloud.

Note
For a Firewall Management Center high-availability pair, you also need to integrate the secondary Firewall Management Center with Cisco Security Cloud.

a) Choose **Integration** > **Cisco Security Cloud**.
b) Click **Enable Cisco Security Cloud** to open a separate browser tab to log you into your Cisco Security Cloud account and confirm the displayed code.

   Make sure this page is not blocked by a pop-up blocker. If you do not already have a Cisco Security Cloud and Security Cloud Control account, you can add one during this procedure.

For detailed information about this integration, see the "System Configuration" chapter in the Cisco Secure Firewall Management Center Administration Guide.

Security Cloud Control onboards the on-prem Firewall Management Center after you integrate the Firewall Management Center with Cisco Security Cloud. Security Cloud Control needs the Firewall Management Center in its inventory for zero-touch provisioning to operate. However, you do not need to use Security Cloud Control directly. If you do use Security Cloud Control, its Firewall Management Center support is limited to device onboarding, viewing its managed devices, viewing objects associated with the Firewall Management Center, and cross-launching the Firewall Management Center.

c) Make sure **Enable Zero-Touch Provisioning** is checked.
d) Click **Save**.

Step 2 Choose **Devices** > **Device Management**.

Step 3 From the **Add** drop-down menu, choose **Device (Wizard)**.

Step 4 Click **Use Serial Number**, and then click **Next**.

Figure 25: Device Registration Method

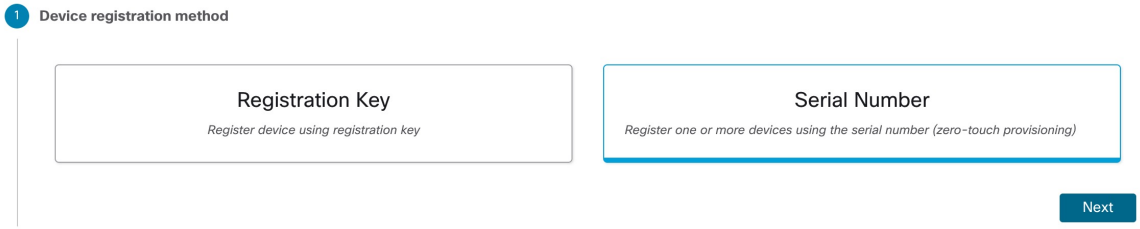

Step 5 In a multi-domain environment, choose the **Domain** from the drop-down list and click **Next**.

Figure 26: Domain

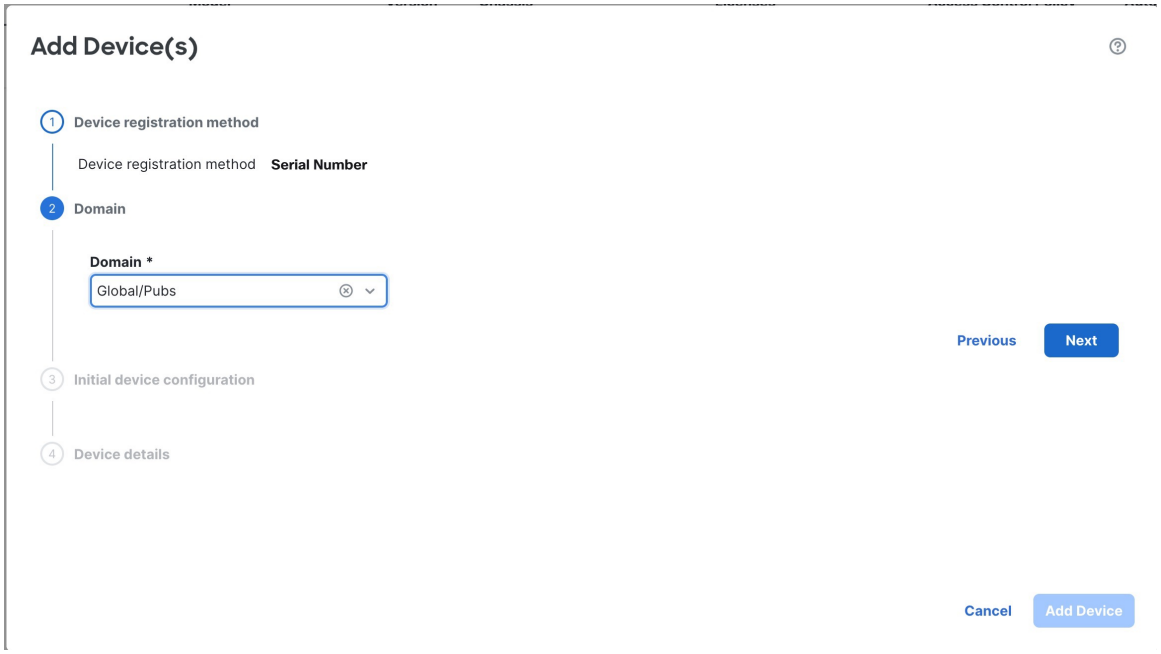

Step 6 For the **Initial device configuration**, click the **Device template** radio button.

Figure 27: Initial Device Configuration

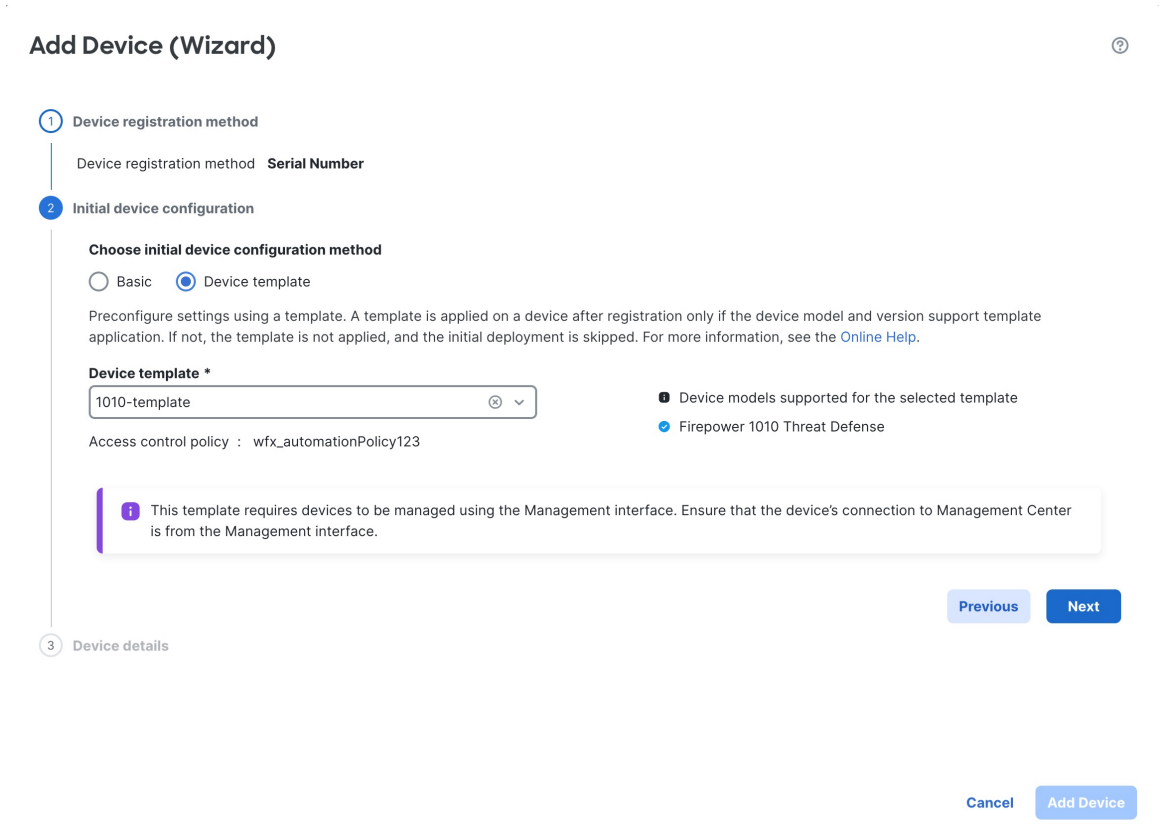

Step 7 Choose the **Device template** from the drop-down list, and click **Next**.

Step 8 In **Device details**, upload a CSV file with the device details required by the template.

Figure 28: Device Details

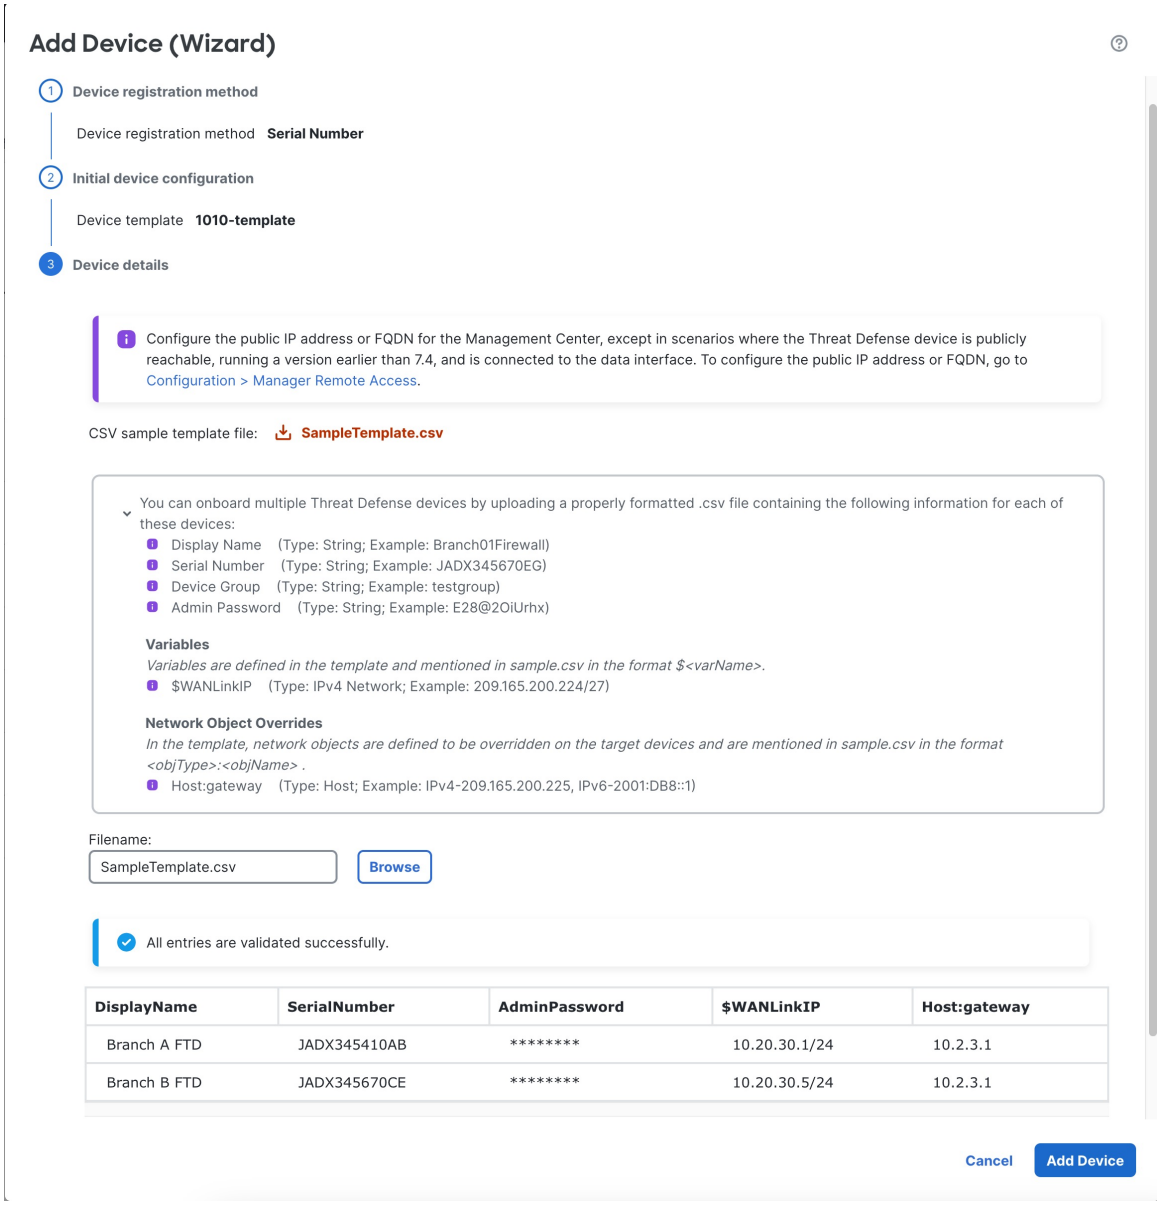

a) Download **SampleTemplate.csv**. This file includes all required headers for values that you need to define per device. For more information on the CSV template file fields, see CSV Template File.

b) **Drag & drop** your CSV template file or **Browse** to select the CSV template file that you want to upload. A validation check is done on the file after you upload it.

After the CSV template file has been uploaded successfully, the content of the CSV template file is displayed in a table format.

Step 9 Click **Add Device** to register the devices.

When using zero-touch provisioning on the outside interface, Security Cloud Control acts as a DDNS provider and does the following:

- Enables DDNS on outside using the "fmcOnly" method. This method is only supported for zero-touch provisioning devices.
- Maps the outside IP address with the following hostname: *serial-number*.**local**.
- Provides the IP address/hostname mapping to the Firewall Management Center so it can resolve the hostname to the correct IP address.
- Informs the Firewall Management Center if the IP address ever changes, for example, if the DHCP lease renews.

If you use zero-touch provisioning on the Management interface, DDNS is not supported. The Firewall Management Center must be publicly reachable so the device can initiate the management connection.

You can continue to use Security Cloud Control as the DDNS provider, or you can later change the DDNS configuration in the Firewall Management Center to a different method. See Configure Dynamic DNS, on page 920 for more information.

If the device fails to register, see Resolve Serial Number (Zero-Touch Provisioning) Registration Issues, on page 57.

CSV Template File for Serial Number Registration with a Device Template

The CSV template file must be less than 2 MB in size. The filename must satisfy the following criteria:

- Can have a maximum of 64 characters.
- Only alphanumeric characters and special characters such as dash (-), period (.), and underscore (_) are allowed.
- Must not contain any spaces.

See the following sample CSV template file containing configuration for two devices.

```
DisplayName,SerialNumber,AdminPassword,$WANLinkIP,Host:gateway
Branch A FTD,JADX345410AB,C15c05n0rt#,10.20.30.1/24,10.2.3.1
Branch B FTD,JADX345670CE,Admin123!,10.20.30.5/24,10.2.3.1
```

A properly formatted CSV file has the following fields.

Mandatory Fields

- **DisplayName**—Name of the device. Type: string. Example: test1
- **SerialNumber**—Serial number of the device. Type: string, Example: JADX345670EG
- **AdminPassword**—(Might be required) Password for admin access, Type: string, Example: E28@2OiUrhx. If this device is unconfigured or a fresh install, then you need to set a **AdminPassword**. If you already logged in and changed the password, leave this field blank.

Optional Fields

- **DeviceGroup**—Name of the device group, Type: string, Example: testgroup

Variables

Use the following format: **$***varName*.

Sample variable: **$LAN-Devices-IPv4Address**—IPv4 address of the LAN device. Type: string. Example: 10.2.3.4/24.

Network Object Overrides

Use the following format: *objType*:*objName*.

Sample network object override: **Network:LAN-Devices-Network**—IP address of the network of LAN devices. Type: string. Example: 10.2.3.0/24

FQDNs

For serial number registration, DDNS is automatically enabled. If you want to set different values from the default for the **FMC Only** type DDNS, then you can configure the settings in the template. In this case, when you provide the CSV value for the hostaname, be sure to specify it as *serialnumber***.local**.

Add a Chassis

You can add a Firepower 4100/9300 chassis to the Firewall Management Center. The management center and the chassis share a separate management connection using the chassis MGMT interface. The Firewall Management Center offers chassis-level health alerts. For configuration, you still need to use the Secure Firewall Chassis Manager or FXOS CLI.

Note For the Secure Firewall 3100/4200, the chassis is added to the Firewall Management Center as part of the conversion to multi-instance mode. See Convert a Device to Multi-Instance Mode, on page 372. However, if you used the CLI to convert to multi-instance mode (Enable Multi-Instance Mode at the CLI, on page 410), skip to Step 3, on page 54 of this procedure to add the chassis to the management center.

Procedure

Step 1 Connect to the chassis FXOS CLI, either using the console port or SSH.

Step 2 Configure the Firewall Management Center.

create device-manager *manager_name* [**hostname** {*hostname* | *ipv4_address* | *ipv6_address*}] [**nat-id** *nat_id*]

You are prompted for the registration key.

You can enter this command from any scope. This command is accepted immediately without using **commit-buffer**.

- **hostname** {*hostname* | *ipv4_address* | *ipv6_address*}—Specifies either the FQDN or IP address of the Firewall Management Center. At least one of the devices, either the Firewall Management Center or the chassis, must have a reachable IP address to establish the two-way, TLS-1.3-encrypted communication channel between the two devices. If you do not specify a **hostname**, then the chassis must have a reachable IP address or hostname and you must specify the **nat-id**.

- **nat-id** *nat_id*—Specifies a unique, one-time string of your choice that you will also specify on the Firewall Management Center when you register the chassis when one side does not specify a reachable IP address or hostname. It is required if you do not specify a **hostname**, however we recommend that you always set the NAT ID even when you specify a hostname or IP address. The NAT ID must not exceed 37 characters. Valid characters include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-). This ID cannot be used for any other devices registering to the Firewall Management Center.
- **Registration Key:** *reg_key*—You will be prompted for a one-time registration key of your choice that you will also specify on the Firewall Management Center when you register the chassis. The registration key must not exceed 37 characters. Valid characters include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-).

Example:

```
firepower# create device-manager boulder_fmc hostname 10.89.5.35 nat-id 93002
(Valid registration key characters: [a-z],[A-Z],[0-9],[-]. Length: [2-36])
Registration Key: Impala67
```

Step 3 In the Firewall Management Center, add the chassis using the chassis management IP address or hostname.

a) Choose **Devices** > **Device Management**, and then **Add** > **Chassis**.

Figure 29: Add Chassis

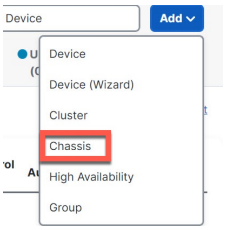

Figure 30: Add Chassis

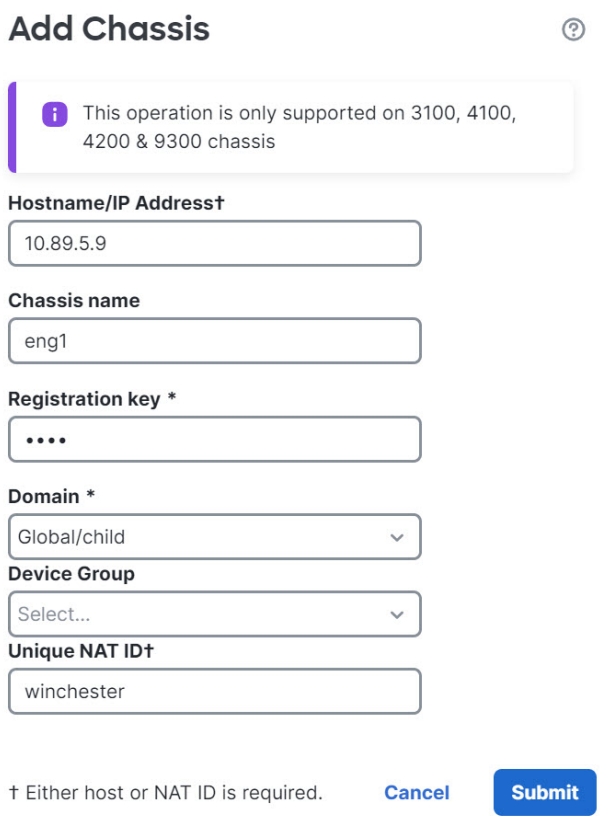

b) In the **Hostname/IP Address** field, enter the IP address or the hostname of the chassis you want to add.

   If you don't know the hostname or IP address, you can leave this field blank specify the **Unique NAT ID**.

c) In the **Chassis Name** field, enter a name for the chassis as you want it to display in the Firewall Management Center.

d) In the **Registration Key** field, enter the same registration key that you used when you configured the chassis to be managed by the Firewall Management Center.

   The registration key is a one-time-use shared secret. The key can include alphanumeric characters and hyphens (-).

e) In a multidomain deployment, regardless of your current domain, assign the chassis to a leaf **Domain**.

   If your current domain is a leaf domain, the chassis is automatically added to the current domain. If your current domain is not a leaf domain, post-registration, you must switch to the leaf domain to configure the chassis. A chassis can only belong to one domain.

f) (Optional) Add the chassis to a **Device Group**.

g) If you used a NAT ID during chassis setup, expand enter the same NAT ID in the **Unique NAT ID** field.

   The NAT ID can include alphanumeric characters and hyphens (-).

h) Click **Submit**.

The chassis is added to the **Devices** > **Device Management** page.

Register With a New Management Center

This procedure shows how to register with a new Firewall Management Center. You should perform these steps even if the new Firewall Management Center uses the old Firewall Management Center's IP address.

Procedure

Step 1 On the old Firewall Management Center, if present, unregister the managed device. See Unregister a Device from the Firewall Management Center, on page 60.

You cannot change the Firewall Management Center IP address if you have an active connection with the Firewall Management Center.

Step 2 Connect to the device CLI, for example using SSH.

Step 3 Configure the new Firewall Management Center.

configure manager add {*hostname* | *IPv4_address* | *IPv6_address* | **DONTRESOLVE** } *regkey* [*nat_id*] [*display_name*]

- {*hostname* | *IPv4_address* | *IPv6_address*}—Sets the Firewall Management Center hostname, IPv4 address, or IPv6 address.
- **DONTRESOLVE**—If the Firewall Management Center is not directly addressable, use **DONTRESOLVE** instead of a hostname or IP address. If you use **DONTRESOLVE** , then a *nat_id* is required. When you add this device to the Firewall Management Center, make sure that you specify both the device IP address and the *nat_id*; one side of the connection needs to specify an IP address, and both sides need to specify the same, unique NAT ID.
- *regkey*—Make up a registration key to be shared between the Firewall Management Center and the device during registration. You can choose any text string for this key between 1 and 37 characters; you will enter the same key on the Firewall Management Center when you add the Firewall Threat Defense.
- *nat_id*—Make up an alphanumeric string from 1 to 37 characters used only during the registration process between the Firewall Management Center and the device when one side does not specify an IP address. This NAT ID is a one-time password used only during registration. Make sure the NAT ID is unique, and not used by any other devices awaiting registration. Specify the same NAT ID on the Firewall Management Center when you add the Firewall Threat Defense.
- *display_name*—Provide a display name for showing this manager with the **show managers** command. This option is useful if you are identifying Security Cloud Control as the primary manager and an on-prem Firewall Management Center for analytics only. If you don't specify this argument, the firewall auto-generates a display name using one of the following methods:
  - *hostname* | *IP_address* (if you don't use the **DONTRESOLVE** keyword)
  - **manager-***timestamp*

Example:

```
> configure manager add DONTRESOLVE abc123 efg456
Manager successfully configured.
Please make note of reg_key as this will be required while adding Device in FMC.

>
```

Step 4 Add the device to the Firewall Management Center.

Resolve Serial Number (Zero-Touch Provisioning) Registration Issues

If the device fails to register using the serial number, the device may not have successfully connected to the cloud. If the **S** or **M** LED is alternating green and amber, this failure might occur because of one of the following causes:

- **Cause**: You performed initial configuration at the CLI and disabled zero-touch provisioning.

  Symptom: Registration will time out.

  Workaround: See Restart Zero-Touch Provisioning at the CLI, on page 59.

- **Cause**: You performed initial configuration in the Firewall Device Manager and disabled zero-touch provisioning.

  Symptom: Registration will time out.

  Workaround: See Restart Zero-Touch Provisioning Using the Firewall Device Manager, on page 59.

- **Cause**: The serial number was already claimed by another manager in your tenant.

  Symptom: You will see a **A device with serial number <serial-number> already exists in the Security Cloud Control tenant** error in the Firewall Management Center.

  Workaround: If you already unregistered the device or you're unsure which Firewall Management Center has claimed it, then see Check for Changes in Security Cloud Control, on page 58.

  If the device is registered to another Firewall Management Center in your tenant for which you don't have access, you can either Reset the Device, on page 60 or Restart Zero-Touch Provisioning at the CLI, on page 59.

- **Cause**: The serial number was already claimed by another manager in another tenant.

  Symptom: You will see a **A device with serial number <serial-number> already exists in another Security Cloud Control tenant** error in the Firewall Management Center.

  Workaround: See either Reset the Device, on page 60 or Restart Zero-Touch Provisioning at the CLI, on page 59.

If you want to bypass any serial number troubleshooting, you can always Use Manual Registration and a Registration Key, on page 60.

For other requirements for serial number registration, see Add a Device Using the Serial Number (Zero-Touch Provisioning)—Basic Configuration, on page 41.

To work around a registration failure, do one of the following tasks.

Check for Changes in Security Cloud Control

If you unregister the device from a Firewall Management Center, Security Cloud Control may not have removed the device from your inventory yet. You can force a refresh in Security Cloud Control to free up the serial number.

1. Log into Security Cloud Control at https://security.cisco.com.
2. Choose **Security Devices**, and then search for your serial number.

   Figure 31: Search by Serial Number

   If the device did not appear in the search results, it may be registered to a different tenant. In this case, see Reset the Device, on page 60 or Restart Zero-Touch Provisioning at the CLI, on page 59.

3. Select the device and click **Check for Changes** on the right.

   Figure 32: Check for Changes

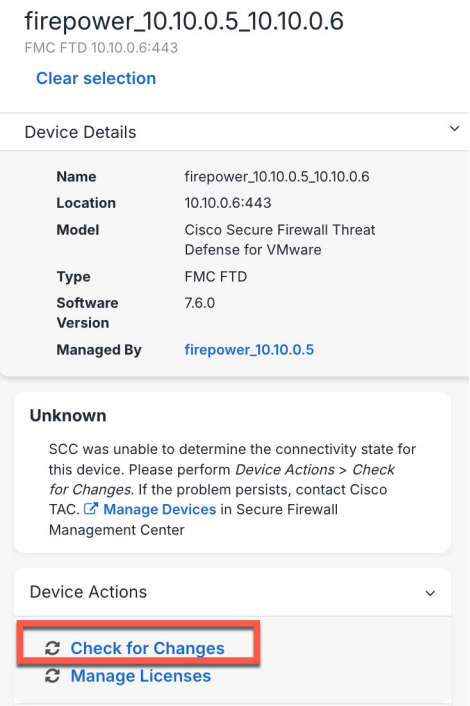

   If the device was unregistered, it will disappear from the list and the device will attempt to connect to the cloud again.

   If the device does not disappear, that means it's registered to another manager in the tenant. The name of the device includes the manager it is registered to, for example, **firepower_10.89.5.36**_1010-1. If you

still want to register the device to a different manager, and you don't have access to the Firewall Management Center to unregister it, you can either Reset the Device, on page 60 or Restart Zero-Touch Provisioning at the CLI, on page 59.

Restart Zero-Touch Provisioning at the CLI

You can unregister the serial number, clear the configuration and any existing management connection, and start the process over.

1. Connect to the FXOS CLI using SSH or the console port.

   If you used SSH, you connect to the Firewall Threat Defense CLI. In this case, enter **connect fxos**. If you used the console port, you connect directly to FXOS.

   ```
   > connect fxos
   firepower#
   ```

2. Enter local management.

   connect local-mgmt

   ```
   firepower# connect local-mgmt
   firepower(local-mgmt)#
   ```

3. Deregister the device from the Cisco cloud.

   cloud deregister

   ```
   firepower(local-mgmt)# cloud deregister
   Release Image Detected RESULT=success MESSAGE=SUCCESS 10, X-Flow-Id:
   2b3c9e8b-76c3-4764-91e4-cfd9828e73f9
   ```

4. Erase the configuration to restore cloud connectivity.

   erase configuration

   ```
   firepower(local-mgmt)# erase configuration
   All configurations will be erased and system will reboot. Are you sure? (yes/no):yes
   Removing all the configuration. Please wait....
   Configurations are cleaned up. Rebooting....
   ```

5. Add a Device Using the Serial Number (Zero-Touch Provisioning)—Basic Configuration, on page 41

Restart Zero-Touch Provisioning Using the Firewall Device Manager

You can accidentally disable low-touch provisioning if you log into the Firewall Device Manager. In this case, you can restart zero-touch provisioning within the Firewall Device Manager.

Note If the serial number was already claimed, see Restart Zero-Touch Provisioning at the CLI, on page 59 instead.

1. In the Firewall Device Manager, click **Device**, then click the **System Settings** > **Cloud Services**.

2. Check **Auto-enroll with** Security Cloud Control Firewall Management **or Secure Firewall Management Center**.
3. Click **Register**.
4. Add a Device Using the Serial Number (Zero-Touch Provisioning)—Basic Configuration, on page 41

Reset the Device

If you do not have access to the CLI and want to make sure your device is unconfigured and ready for zero-touch provisioning, reset the device to its default state by press the small, recessed Reset button for longer than five seconds. See your hardware installation guide for more information.

Use Manual Registration and a Registration Key

If zero-touch provisioning fails, the easiest way to complete registration is to use the registration key method.

1. See Complete the Firewall Threat Defense Initial Configuration for Manual Registration, on page 13 or Complete the Firewall Threat Defense Initial Configuration Using the Firewall Device Manager, on page 13.
2. If you are not presented with the initial setup tasks, it's possible your device was successfully registered to another Firewall Management Center. You must first delete the management connection and then re-register with the correct manager.
   a. First, check if registration has completed:

```
> show managers
Type                    : Manager
Host                    : 10.10.1.4
Display name            : 10.10.1.4
Identifier              : f7ffad78-bf16-11ec-a737-baa2f76ef602
Registration            : Completed
Management type         : Configuration
```

   b. If **Registration** shows **Completed**, you need to delete the manager:

      configure manager delete

   c. You can then register the device at the CLI using **configure manager add**.

Unregister a Device from the Firewall Management Center

If you no longer want to manage a device, you can unregister it from the Firewall Management Center.

To unregister a cluster, cluster node, or high availability pair, see the chapters for those deployments.

Unregistering a device:

- Severs all communication between the Firewall Management Center and the device.
- Removes the device from the **Device Management** page.
- Returns the device to local time management if the device's platform settings policy is configured to receive time from the Firewall Management Center using NTP.

- Leaves the configuration intact, so the device continues to process traffic.

  Policies, such as NAT and VPN, ACLs, and the interface configurations remain intact.

Registering the device again to the same or a different Firewall Management Center causes the configuration to be removed, so the device will stop processing traffic at that point.

Before you unregister the device, be sure to export the configuration or create a template so you can re-apply the device-level configuration (interfaces, routing, and so on) when you re-register it. If you do not have a saved configuration or template, you will have to re-configure device settings.

After you re-add the device and either import a saved configuration, use a template, or re-configure your settings, you need to deploy the configuration before it starts passing traffic again.

Before you begin

To re-apply the device-level configuration if you re-add it to the Firewall Management Center, do one of the following:

- Export the device configuration. See Export and Import the Device Configuration, on page 136.
- Create a template for the device.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to unregister, click **More** (⋮), and then click **Unregister**.

Figure 33: Unregister

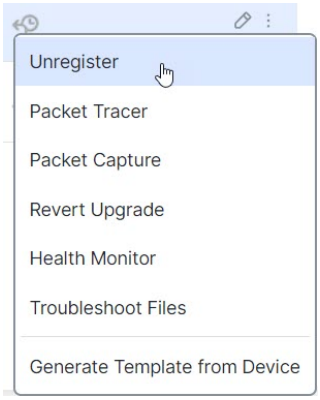

Step 3 Confirm that you want to unregister the device.

Step 4 You can now change your manager.

- Re-register the device to this Firewall Management Center—If you know the registration key and NAT ID, refer to Registration Key Method, on page 31. If you need to reset them, you can reconfigure the manager as though it's new. See Register With a New Management Center, on page 56.
- Register to a new Firewall Management Center—Register With a New Management Center, on page 56.

- Change to the Firewall Device Manager—Switch from Firewall Management Center to Firewall Device Manager, on page 77.
- Delete the manager without specifying a new one—To sever the management connection on the Firewall Threat Defense without identifying a new manager (no manager mode), from the Firewall Threat Defense CLI use the **configure manager delete** command.

Shut Down or Restart the Device

It's important that you shut down your system properly. Simply unplugging the power or pressing the power switch can cause serious file system damage. Remember that there are many processes running in the background all the time, and unplugging or shutting off the power does not allow the graceful shutdown of your firewall.

See the following task to shut down or restart your system properly.

Note After restarting your device, you may see an error that the management connection could not be reestablished. In some cases, the connection is attempted before the Management interface on the device is ready. The connection will be retried automatically and should come up within 15 minutes.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device that you want to restart, click **Edit** ( ).

Step 3 Click **Device**.

Step 4 To restart the device:

a) Click **Restart Device** ( ).
b) When prompted, confirm that you want to restart the device.

Step 5 To shut down the device:

a) Click **Shut Down Device** ( ) in the **System** section.
b) When prompted, confirm that you want to shut down the device.
c) If you have a console connection to the firewall, monitor the system prompts as the firewall shuts down. You will see the following prompt:

```
System is stopped.
It is safe to power off now.
Do you want to reboot instead? [y/N]
```

If you do not have a console connection, wait approximately 3 minutes to ensure the system has shut down.

For the ISA 3000, when shutdown is complete, the System LED will turn off. Wait at least 10 seconds before you remove the power.

Download the Managed Device List

You can download a report of all the managed devices.

Before you begin

To perform the following task, you must be an Admin user.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click the **Download Device List Report** link.

Step 3 You can download the device list in CSV or PDF format. Choose **Download CSV** or **Download PDF** to download the report.

Migrate Firewall Threat Defense Devices

The Secure Firewall Threat Defense model migration wizard enables you to migrate configurations from an earlier Firewall Threat Defense model. After the migration, all routing and interface configurations from the source Firewall Threat Defense device are available in the target Firewall Threat Defense.

The wizard supports multiple models as source and target devices, for more information see Supported Devices for Migration, on page 63.

Supported Devices for Migration

Supported Source Devices

- Cisco Firepower 1120
- Cisco Firepower 1140
- Cisco Firepower 1150
- Cisco Firepower 2110
- Cisco Firepower 2120
- Cisco Firepower 2130
- Cisco Firepower 2140
- Cisco Firepower 4110
- Cisco Firepower 4120

- Cisco Firepower 4140
- Cisco Firepower 4150
- Cisco Firepower 9300 Series SM-24
- Cisco Firepower 9300 Series SM-36
- Cisco Firepower 9300 Series SM-44

Note The source devices must be Version 7.2.x and later.

Supported Target Devices

- Cisco Secure Firewall 3105
- Cisco Secure Firewall 3110
- Cisco Secure Firewall 3120
- Cisco Secure Firewall 3130
- Cisco Secure Firewall 3140
- Cisco Firepower 4215
- Cisco Firepower 4225
- Cisco Firepower 4245

Note The target devices must be Version 7.4.1 and later.

Supported Migration Paths

The following table lists the supported target Firewall Threat Defense models that you can migrate to from your source Firewall Threat Defense model.

Source Model	Target Model			
	Cisco Secure Firewall 3100 Series	Cisco Secure Firewall 4200 Series	Instance in Secure Firewall 3100 Series	Instance in Secure Firewall 4200 Series
Firepower 1100 Series	Yes	—	—	—
Firepower 2100 Series	Yes	—	—	—
Firepower 4100 Series	Yes	Yes	—	—
Firepower 9300 Series	Yes	Yes	—	—
Instance from Firepower 4100 Series	—	—	Yes	Yes

Instance from Firepower 9300 Series	—	—	Yes	Yes

License for Migration

- Your Smart Licensing account must have the license entitlements for the target device.
- You must register and enroll the device with the Smart Licensing account. The migration copies the source device licenses to the target device.

Prerequisites for Migration

- **General device prerequisites**
  - Register the source and the target devices to the Firewall Management Center.
  - Ensure that the target device is a newly registered device without any configurations.
  - Source and target devices must be in the same state and modes:
    - Domain
    - Firewall mode: Routed or Transparent
    - Compliance mode (CC or UCAPL)
    - Management state

      Devices must have the same type of manager access interfaces (management interface or data interface).
    - Multi-instance mode or appliance mode
  - Ensure that you have permission for modifications on the devices.
  - Ensure that the configurations on the source device are valid and have no errors.
  - Deployment, import, or export tasks must not run on either of the devices during the migration. The source device can have pending deployments.
- **Prerequisites for change management**
  - Ensure that source and target devices are not locked by a change management ticket.
  - Ensure that shared policies assigned to the source device are not locked by a change management ticket.
- **Prerequisites for HA devices**
  - Migrate a device only from an active Firewall Management Center.
- **Prerequisites for devices in multi-instance mode**
  - Ensure that the source and target devices are in multi-instance mode.
  - Manually migrate the chassis configurations. Create instances before migrating the instance configuration to the target instances. The target device must have compatible interfaces. For example,

on the target device, you must create EtherChannel interfaces, and also create tagged, untagged, dedicated, or shared interfaces for these interfaces on the target device.

- **Prerequisite for devices with out-of-band configurations**
  - Ensure that you acknowledge out-of-band changes and match the configurations within the Firewall Management Center. You cannot migrate devices with these configurations. To view out-of-band configurations:
    1. Choose **Devices** > **Device Management**.
    2. Click the edit icon next to the device and click the **Interfaces** tab.

- **Prerequisites for devices with manager access interfaces**

  Ensure that the devices are not in Data Transit or Management Transit states. You cannot migrate if devices are in these states.
  - Data Transit state: Device state when the manager access interface changes from data interface to management interface without deploying the changes on the device.
  - Management Transit state: Device state when the manager access interface changes from management interface to data interface without deploying the changes on the device.

- **Prerequisite for devices with merged management and diagnostic interfaces**

  Ensure that the target device is always in merged mode.

What Configurations Does the Wizard Migrate?

The migration wizard copies the following configurations from the source device to the target device:

- Licenses
- Interface configurations
- Inline sets configurations
- Routing configurations
- DHCP and DDNS configurations
- Virtual router configurations
- Policies
- Associated objects and object overrides
- Platform settings
- Remote branch deployment configurations

The migration wizard copies the following policy configurations from the source device to the target device:

- Health policy
- NAT policy

- QoS policy
- Remote access VPN policy
- FlexConfig policy
- Access control policy
- Prefilter policy
- IPS policy
- DNS policy
- SSL policy
- Malware and File policy
- Identity policy
- Shared policy

The migration wizard copies the following routing configurations from the source device to the target device:

- ECMP
- BFD
- OSPFv2/v3
- EIGRP
- RIP
- BGP
- Policy Based Routing
- Static Route
- Multicast Routing
- Virtual Router

The migration wizard copies the following interfaces from the source device to the target device:

- Physical interfaces
- Sub-interfaces
- EtherChannel interfaces
  - On a standalone device, the wizard copies the EtherChannels from the source device to the target device.
  - For devices in multi-instance mode, you must create EtherChannels on the chassis and assign them to the instance.
- Bridge group interfaces
- VTI interfaces

- VNI interfaces
- Loopback interfaces
- Inline interfaces
- VXLAN tunnel endpoint (VTEP) interfaces

The migration wizard retains the device group of the target device.

Guidelines and Limitations for Migration

Guidelines

- **For devices in multi-instance mode**:

  During migration, ensure that you map the interfaces according to the table below:

Source Device	Target Device
Physical interface	Physical interface
EtherChannel interface	EtherChannel interface
Supervisor-provisioned subinterface	Supervisor-provisioned subinterface
Tagged interface	Tagged interface
Untagged interface	Untagged interface
Shared interface	Shared and dedicated interface
Dedicated interface	Dedicated interface

  You cannot map a supervisor-provisioned subinterface to a subinterface created by an instance.

- **For HA devices**, you can migrate:
  - Source HA device to target HA device.
  - Source HA device to target standalone device.
- **For devices in remote branch deployment**:
  - Map the source manager access interface to the target manager access interface.
  - Ensure that the manager access interfaces of the source and target Firewall Management Centers are of the same IP address type (static or DHCP).
  - Both manager access interfaces must have IPv4 or IPv6 addresses.
  - If the manager access interfaces have static IP addresses, ensure that they are in the same subnet.
- **For Snort**:
  - If target device has Snort 3, after migration it will have Snort 3.
  - If the source and target devices have Snort 2, after migration the target device will have Snort 2.

- **For devices using diagnostic interfaces**:

  Only merged management interfaces are available on the target devices after migration.

Limitations

- The migration wizard does not migrate:
  - Site-to-site VPN policies
  - SNMP device configurations for Firepower 2100 Series

    After the migration, you can configure SNMP using the platform settings for the device.
- You can perform only one migration at a time.
- Remote access VPN trustpoint certificates are not enrolled after migration.
- For HA devices:
  - Target device: You cannot migrate a standalone device to an HA device.
- Clustering is not supported.
- For devices in remote branch deployment:
  - The wizard does not migrate a single WAN manager access data interface to a dual WAN manager access data interface.

Migrate a Secure Firewall Threat Defense

Before you begin

Ensure you review Prerequisites for Migration, on page 65 and Guidelines and Limitations for Migration, on page 68.

Procedure

Step 1 Choose **Firewall Devices > Device Management**.

Step 2 Click **Migrate** in the top right corner of the page.

Step 3 In **Select source and target devices**:

a) From the **Source device** drop-down list, choose a device.

b) From the **Target device** drop-down list, choose a device.

The source and target devices can have these tags:

- Routed: Devices in routed firewall mode.
- Transparent: Devices in transparent firewall mode.
- Container: Devices in multi-instance mode.
- High Availability: Devices in high availability mode.

- Analytics Only: Devices managed by Security Cloud Control and the Firewall Management Center only recieves and displays the events (analytics-only Firewall Management Center).

If the device is part of an HA pair, only the HA pair name appears.

Step 4 Click **Next**.

Step 5 (Only for Firepower 4100 and 9300 Series devices in appliance mode) In **Chassis manager details**:

a) Check the **Skip chassis manager** check box, if required.

b) In the **Chassis hostname or IP address** field, enter the values.

Note

- Verify that the Secure Firewall Chassis Manager is reachable from the Firewall Management Center.
- Ensure you select the correct chassis manager for the source device, as Firewall Management Center does not validate your choice.

c) Click **Verify certificate** to verify the chassis manager's certificate.

d) In the **Username** and **Password** fields, enter the credentials of the chassis manager.

Step 6 Click **Next**.

Step 7 In **Configure interfaces**:

By default, the source and target interfaces are mapped using the interface hardware name. You must map named interfaces, logical interfaces, and interfaces that are part of other interfaces. Mapping of all other interfaces is not mandatory. The wizard creates the logical interfaces according to the interface mapping that you provide.

You cannot map interfaces that are part of an HA failover configuration. These interfaces are disabled in the wizard.

Only data interfaces are available for interface mapping. Management, eventing, and diagnostic interfaces are not available for the interface mapping.

Firepower 4100 and 9300 Series devices in appliance mode:

For these devices, the Firewall Management Center fetches interface attributes such as speed, duplex, and auto-negotiation from the chassis manager.

a) Click one of the following options to configure these interface attributes on the target device:

- **Retain target device values**: (Default) Retains the interface attributes configured on the target device.
- **Copy from source device**: Copies the interface attributes from the source device.

  This option is enabled only when Firewall Management Center successfully connects to the chassis manager. We recommend that you use this option. The speed, duplex, and auto-negotiation values of physical interfaces are set to default values if they are incompatible in the target device.
- **Customize device values**—Allows you to configure the values of the required interface attributes on the target device.

b) To change the interface mapping from the default ones, choose an interface from the **Mapped interface** drop-down list.

c) For EtherChannels, you can configure interface attributes and click **Add member interface** to add member interfaces.

Interface attributes of an EtherChannel is configured based on the first member interface's interface attributes.You can add up to 16 member interfaces.

Firepower 1100 and 2100 Series devices, and Firepower 4100 and 9300 Series devices in multi-instance mode:

For these devices, you must map the source device interfaces to target device interfaces.

For Firepower 4100 and 9300 Series devices in multi-instance mode, you can only perform the interface mapping and you cannot configure the interface attributes such as speed, duplex, auto-negotiation, and FEC mode.

If you want to change the interface mapping from the default ones, choose an interface from the **Mapped interface** drop-down list.

Click **Reset** to configure the default interface mappings. For example, the wizard maps Ethernet1/1 in the source device to Ethernet1/1 in the target device.

The interfaces can have the following tags:

- Tagged: Physical interfaces on the chassis.
- Untagged: Physical interfaces on the chassis that have sub-interfaces.
- Dedicated: Interfaces that are assigned to specific instances and are not shared across multiple instances.
- Shared: Interfaces that are shared by multiple instances.
- Manager access: Data interface is the manager access interface.

Check the **Ignore warning** check box, if required.

Step 8 Click **Next**.

Step 9 Click **Submit** to start the migration.

Step 10 View the migration status on the **Notifications** > **Tasks** page.

A **Device Model Migration** report is generated after the migration is completed. You will see a link to this report in the **Notifications** > **Tasks** page.

What to do next

After a successful migration, you must complete these tasks:

- Review the recommendations in Best Practices for Threat Defense Device Migration, on page 71.
- Validate the configurations.
- Deploy the configurations on the device.

In case of a migration failure, the target device is rolled back to the initial state.

Best Practices for Threat Defense Device Migration

After a successful migration, we recommend that you perform the following actions before the deployment:

- IP addresses of the interfaces are copied to the target device from the source device. Change the IP addresses of the target device interfaces, if the source device is live
- Ensure that you update your NAT policies with the modified IP addresses.
- Configure the interface speeds if they are set to default values after migration.
- Re-enroll the device certificates, if any, on the target device.
- (Optional) Configure SNMP for Firepower 1100 and 2100 using the platform settings for the device.
- (Optional) Configure remote branch deployment configurations.

  If the source or target device had manager access through a data interface, after the migration, the manager access will be lost. Update the manager access configuration on the target device. For more information, see the *Change the Manager Access Interface from Management to Data* topic in the Cisco Secure Firewall Management Center Device Configuration Guide or the Online Help.
- Configure site-to-site VPN, if required. These configurations are not migrated from the source device.
- View the deployment preview before the deployment. Choose **Deploy > Advanced Deploy** and click the **Preview** ( ) icon for the device.
- Monitor the health of the device in the health monitor (choose **Troubleshooting > Health > Monitor**). After migration, the health policy of the source device becomes the health policy of the target device. You can also configure a new health policy for the device.

  After migration, the device monitoring dashboard may temporarily display redundant colored lines because the device has different UUIDs before and after migration. This redundancy appears only during the migration time. An hour after migration, the dashboard will show a single line per metric.

Switch Managers

You can change between managers if needed.

Switch from the Firewall Device Manager to the Firewall Management Center

When you switch from the Firewall Device Manager to the Firewall Management Center, all interface configuration is retained, in addition to the Management interface and the manager access settings. Note that other configuration settings, such as the access control policy or security zones, are not retained.

After you switch to the Firewall Management Center, you can no longer use the Firewall Device Manager to manage the Firewall Threat Defense device.

Before you begin

If the firewall is configured for high availability, you must first break the high availability configuration using the Firewall Device Manager (if possible) or the **configure high-availability disable** command. Ideally, break high availability from the active unit.

Procedure

Step 1 In the Firewall Device Manager, unregister the device from the Cisco Smart Software Manager.

Step 2 (Might be required) Configure the Management interface.

You may need to change the Management interface configuration, even if you intend to use a data interface for manager access. You will have to reconnect to the Firewall Device Manager if you were using the Management interface for the Firewall Device Manager connection.

- Data interface for manager access—The Management interface must have the gateway set to data interfaces. By default, the Management interface receives an IP address and gateway from DHCP. If you do not receive a gateway from DHCP (for example, you did not connect this interface to a network), then the gateway will default to data interfaces, and you do not need to configure anything. If you did receive a gateway from DHCP, then you need to instead configure this interface with a static IP address and set the gateway to data interfaces.
- Management interface for manager access—If you want to configure a static IP address, be sure to also set the default gateway to be a unique gateway instead of the data interfaces. If you use DHCP, then you do not need to configure anything assuming you successfully get the gateway from DHCP.

Step 3 Choose **Device** > **System Settings** > **Central Management**, and click **Proceed** to set up the Firewall Management Center management.

Step 4 Configure the **Management Center/SCC Details**.

Figure 34: Management Center/SCC Details

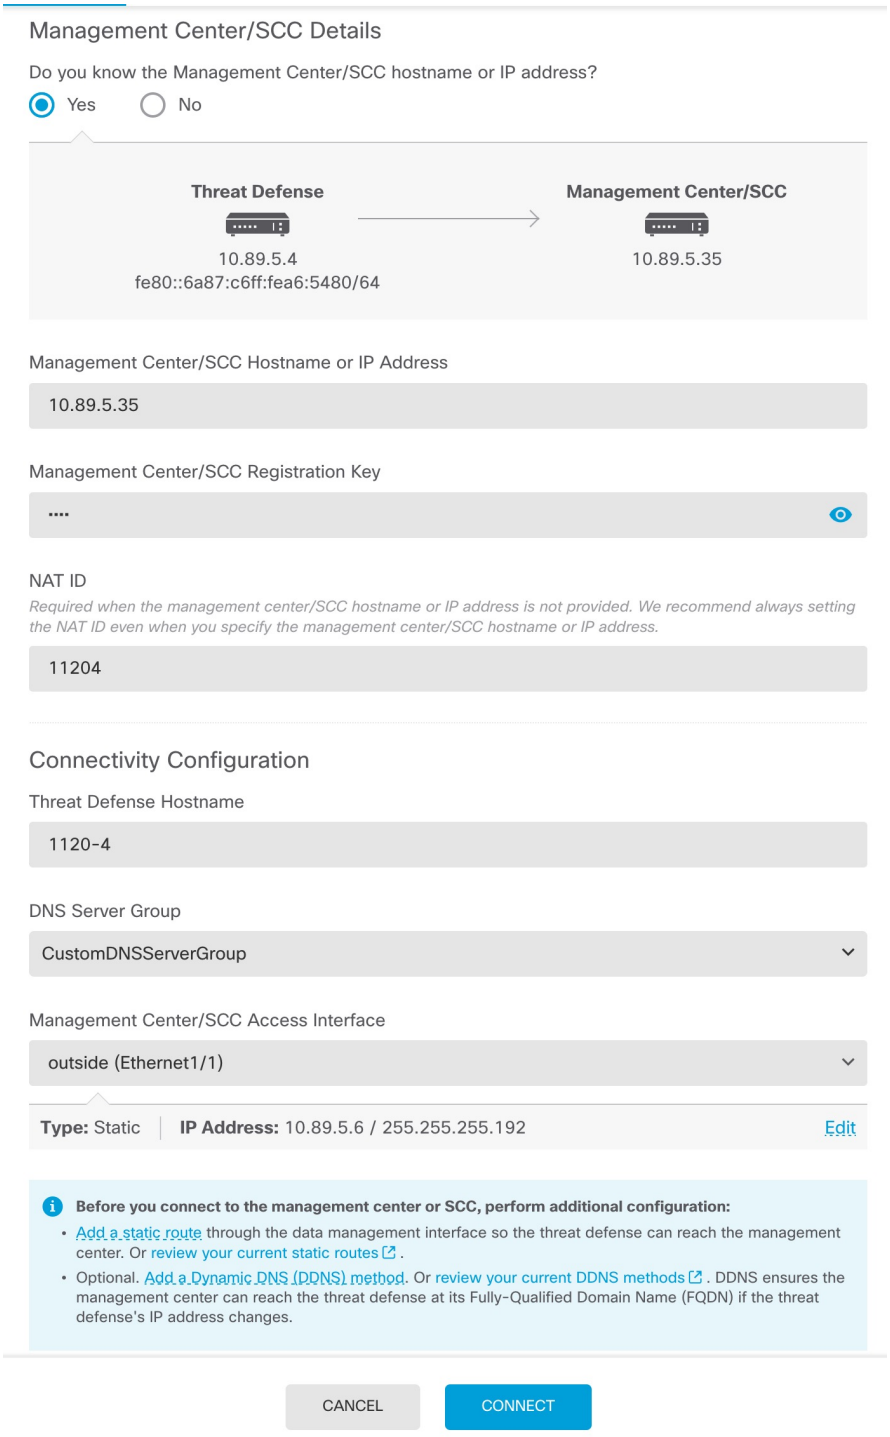

a) For **Do you know the Management Center/SCC hostname or IP address?**, click **Yes** if you can reach the Firewall Management Center using an IP address or hostname, or **No** if the Firewall Management CenterSecurity Cloud Control is behind NAT or does not have a public IP address or hostname.

At least one of the devices, either the Firewall Management Center or the Firewall Threat Defense device, must have a reachable IP address to establish the two-way, TLS-1.3-encrypted communication channel between the two devices.

b) If you chose **Yes**, then enter the **Management Center/SCC Hostname or IP Address**.

c) Specify the **Management Center/SCC Registration Key**.

This key is a one-time registration key of your choice that you will also specify on the Firewall Management Center when you register the Firewall Threat Defense device. The registration key must be between 2 and 36 characters. Valid characters include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-). This ID can be used for multiple devices registering to the Firewall Management Center.

a) Specify a **NAT ID**.

This ID is a unique, one-time string of your choice that you will also specify on the Firewall Management Center. The NAT ID must be between 2 and 36 characters. Valid characters include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-). This ID *cannot* be used for any other devices registering to the Firewall Management Center. The NAT ID is used in combination with the IP address to verify that the connection is coming from the correct device; only after authentication of the IP address/NAT ID will the registration key be checked. We recommended that you always use the NAT ID even when it is optional, but it is required if:

- You set the Firewall Management Center IP address to **DONTRESOLVE**.
- When adding the device on the Firewall Management Center, you do not specify a reachable device IP address or hostname.
- You use the data interface for management, even if you specify IP addresses on both sides.
- The Firewall Management Center uses multiple management interfaces.

Step 5 Configure the **Connectivity Configuration**.

a) Specify the **FTD Hostname**.

If you use a data interface for the **Management Center/SCC Access Interface** access, then this FQDN will be used for this interface.

b) Specify the **DNS Server Group**.

Choose an existing group, or create a new one. The default DNS group is called **CiscoUmbrellaDNSServerGroup**, which includes the OpenDNS servers.

If you intend to choose a data interface for the **Management Center/SCC Access Interface**, then this setting sets the *data* interface DNS server. The Management DNS server that you set with the setup wizard is used for management traffic. The data DNS server is used for DDNS (if configured) or for security policies applied to this interface. You are likely to choose the same DNS server group that you used for Management, because both management and data traffic reach the DNS server through the outside interface.

On the Firewall Management Center, the data interface DNS servers are configured in the Platform Settings policy that you assign to this Firewall Threat Defense device. When you add the Firewall Threat Defense device to the Firewall Management Center, the local setting is maintained, and the DNS servers are *not* added to a Platform Settings policy. However, if you later assign a Platform Settings policy to the Firewall Threat Defense device that includes a DNS configuration, then that configuration will overwrite the local setting. We suggest that you actively configure the DNS Platform Settings to match this setting to bring the Firewall Management Center and the Firewall Threat Defense device into sync.

Also, local DNS servers are only retained by the Firewall Management Center if the DNS servers were discovered at initial registration.

If you intend to choose the Management interface for the **Management Center/SCC Access InterfaceFMC Access Interface**, then this setting configures the Management DNS server.

c) For the **Management Center/SCC Access Interface**, choose any configured interface.

You can change the manager interface after you register the Firewall Threat Defense device to the Firewall Management Center, to either the Management interface or another data interface.

Step 6 (Optional) If you chose a data interface, and it was not the outside interface, then add a default route.

You will see a message telling you to check that you have a default route through the interface. If you chose outside, you already configured this route as part of the setup wizard. If you chose a different interface, then you need to manually configure a default route before you connect to the Firewall Management Center.

If you chose the Management interface, then you need to configure the gateway to be a unique gateway before you can proceed on this screen.

Step 7 (Optional) If you chose a data interface, click **Add a Dynamic DNS (DDNS) method**.

DDNS ensures the Firewall Management Center can reach the Firewall Threat Defense device at its Fully-Qualified Domain Name (FQDN) if the IP address changes. See **Device** > **System Settings** > **DDNS Service** to configure DDNS.

If you configure DDNS before you add the Firewall Threat Defense device to the Firewall Management Center, the Firewall Threat Defense device automatically adds certificates for all of the major CAs from the Cisco Trusted Root CA bundle so that the Firewall Threat Defense device can validate the DDNS server certificate for the HTTPS connection. Firewall Threat Defense supports any DDNS server that uses the DynDNS Remote API specification (https://help.dyn.com/remote-access-api/).

DDNS is not supported when using the Management interface for manager access.

Step 8 Click **Connect**. The **Registration Status** dialog box shows the current status of the switch to the Firewall Management Center. After the **Saving Management Center/SCC Registration Settings** step, go to the Firewall Management Center, and add the firewall.

If you want to cancel the switch to the Firewall Management Center, click **Cancel Registration**. Otherwise, do not close the Firewall Device Manager browser window until after the **Saving Management Center/SCC Registration Settings** step. If you do, the process will be paused, and will only resume when you reconnect to the Firewall Device Manager.

If you remain connected to the Firewall Device Manager after the **Saving Management Center/SCC Registration Settings** step, you will eventually see the **Successful Connection with Management Center/SCC** dialog box, after which you will be disconnected from the Firewall Device Manager.

Figure 35: Successful Connection

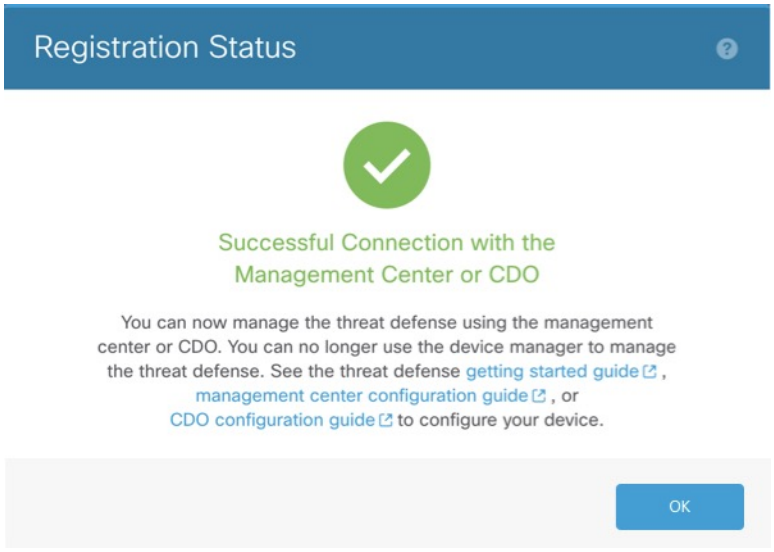

Switch from Firewall Management Center to Firewall Device Manager

You can configure the Firewall Threat Defense device currently being managed by the on-premises or cloud-delivered Firewall Management Center to use the Firewall Device Manager instead.

You can switch from the Firewall Management Center to the Firewall Device Manager without reinstalling the software. Before switching from the Firewall Management Center to the Firewall Device Manager, verify that the Firewall Device Manager meets all of your configuration requirements. If you want to switch from the Firewall Device Manager to the Firewall Management Center, see Switch from the Firewall Device Manager to the Firewall Management Center, on page 72.

Caution Switching to the Firewall Device Manager erases the device configuration and returns the system to the default configuration. However, the Management IP address and hostname are preserved.

Procedure

Step 1 In the Firewall Management Center, unregister the firewall from the **Devices** > **Device Management** page.

Step 2 Connect to the Firewall Threat Defense CLI using SSH or the console port. For SSH, open a connection to the **management IP address**, and log into the Firewall Threat Defense CLI with the **admin** username (or any other user with admin privileges).

The console port defaults to the FXOS CLI. Connect to the Firewall Threat Defense CLI using the **connect ftd** command. The SSH session connects directly to the Firewall Threat Defense CLI.

If you cannot connect to the management IP address, do one of the following:

- Ensure that the Management physical port is wired to a functioning network.

- Ensure that the management IP address and gateway are configured for the management network. Use the **configure network ipv4/ipv6 manual** command.

Step 3 Verify you are currently in remote management mode.

show managers

Example:

```
> show managers
Type                      : Manager
Host                      : 10.89.5.35
Display name              : 10.89.5.35
Identifier                : f7ffad78-bf16-11ec-a737-baa2f76ef602
Registration              : Completed
```

Step 4 Delete the remote manager and go into no manager mode.

configure manager delete *uuid*

You cannot go directly from remote management to local management. If you have more than one manager defined, you need to specify the identifier (also known as the UUID; see the **show managers** command). Delete each manager entry separately.

Example:

```
> configure manager delete
Deleting task list
Manager successfully deleted.

>
> show managers
No managers configured.
```

Step 5 Configure the local manager.

configure manager local

You can now use a web browser to open the local manager at **https://***management-IP-address*.

Example:

```
> configure manager local
Deleting task list

> show managers
Managed locally.
```

Hot Swap an SSD on the Secure Firewall 3100/4200

If you have two SSDs, they form a RAID when you boot up. You can perform the following tasks at the Firewall Threat Defense CLI while the firewall is powered up:

- Hot swap one of the SSDs—If an SSD is faulty, you can replace it. Note that if you only have one SSD, you cannot remove it while the firewall is powered on.
- Remove one of the SSDs—If you have two SSDs, you can remove one.
- Add a second SSD—If you have one SSD, you can add a second SSD and form a RAID.

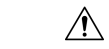

Caution Do not remove an SSD without first removing it from the RAID using this procedure. You can cause data loss.

Procedure

Step 1 Remove one of the SSDs.

a) Remove the SSD from the RAID.

configure raid remove-secure local-disk {**1** | **2**}

The **remove-secure** keyword removes the SSD from the RAID, disables the self-encrypting disk feature, and performs a secure erase of the SSD. If you only want to remove the SSD from the RAID and want to keep the data intact, you can use the **remove** keyword.

Example:

```
> configure raid remove-secure local-disk 2
```

b) Monitor the RAID status until the SSD no longer shows in the inventory.

show raid

After the SSD is removed from the RAID, the **Operability** and **Drive State** will show as **degraded**. The second drive will no longer be listed as a member disk.

Example:

```
> show raid
Virtual Drive
ID:                         1
Size (MB):                  858306
Operability:                operable
Presence:                   equipped
Lifecycle:                  available
Drive State:                optimal
Type:                       raid
Level:                      raid1
Max Disks:                  2
Meta Version:               1.0
Array State:                active
Sync Action:                idle
Sync Completed:             unknown
Degraded:                   0
Sync Speed:                 none

RAID member Disk:
Device Name:                nvme0n1
```

```
Disk State:                 in-sync
Disk Slot:                  1
Read Errors:                0
Recovery Start:             none
Bad Blocks:
Unacknowledged Bad Blocks:

Device Name:                nvme1n1
Disk State:                 in-sync
Disk Slot:                  2
Read Errors:                0
Recovery Start:             none
Bad Blocks:
Unacknowledged Bad Blocks:

> show raid
Virtual Drive
ID:                         1
Size (MB):                  858306
Operability:                degraded
Presence:                   equipped
Lifecycle:                  available
Drive State:                degraded
Type:                       raid
Level:                      raid1
Max Disks:                  2
Meta Version:               1.0
Array State:                active
Sync Action:                idle
Sync Completed:             unknown
Degraded:                   1
Sync Speed:                 none

RAID member Disk:
Device Name:                nvme0n1
Disk State:                 in-sync
Disk Slot:                  1
Read Errors:                0
Recovery Start:             none
Bad Blocks:
Unacknowledged Bad Blocks:
```

c) Physically remove the SSD from the chassis.

Step 2 Add an SSD.

a) Physically add the SSD to the empty slot.

b) Add the SSD to the RAID.

configure raid add local-disk {**1** | **2**}

It can take several hours to complete syncing the new SSD to the RAID, during which the firewall is completely operational. You can even reboot, and the sync will continue after it powers up. Use the **show raid** command to show the status.

If you install an SSD that was previously used on another system, and is still locked, enter the following command:

configure raid add local-disk {**1** | **2**} *psid*

The *psid* is printed on the label attached to the back of the SSD. Alternatively, you can reboot the system, and the SSD will be reformatted and added to the RAID.

Disable the USB Port

By default, the type-A USB port is enabled. You might want to disable USB port access for security purposes. Disabling USB is supported on the following models:

- Firepower 1000 Series
- Secure Firewall 3100
- Secure Firewall 4200

Guidelines

- Enabling or disabling the USB port requires a reboot.
- If the USB port is disabled and you downgrade to a version that does not support this feature, the port will remain disabled, and you cannot re-enable it without erasing the NVRAM (the FXOS local-mgmt **erase secure all** command).
- If you perform a ROMMON **factory-reset** or FXOS local-mgmt **erase secure**, the USB port will be re-enabled.
- For high availability or clustering, you must disable or re-enable the port individually on each unit.

Note This feature does not affect the USB console port, if present.

Disable the USB Port on a Device

To disable the USB port on a device, you can do so at the Firewall Threat Defense CLI.

Procedure

Step 1 Disable the USB port.

system support usb configure disable

reboot

To re-enable the USB port, enter **system support usb configure enable**.

Example:

```
>system support usb configure disable
USB Port Admin State set to 'disabled'.
Please reboot the system to apply any control state changes.
```

```
>reboot
This command will reboot the system. Continue?
Please enter 'YES' or 'NO': YES
```

Step 2 View the port status.

system support usb show

The Admin State shows the USB port configuration. The Oper State shows the current operation. For example, if you disable the USB port but do not reload, the Admin State will show disabled while the Oper State would will enabled.

Example:

```
>system support usb show
USB Port Info
---------------
Admin State: disabled
Oper State: disabled
```

Disable the USB Port in Multi-Instance Mode

To disable the USB port in multi-instance mode, you can do so at the FXOS CLI.

Procedure

Step 1 Disable the USB port and reboot for the change to take effect.

a) Disable the USB port.

scope fabric-interconnect

disable usb-port

commit buffer

b) Reboot the chassis.

connect local-mgmt

reboot

Example:

```
firepower-4245 /fabric-interconnect # disable usb-port
Note: USB enablement or disablement changes are effected only after FXOS reboot.
Confirm change? (yes/no) [yes]:
device /fabric-interconnect* # commit buffer
Note: USB enablement or disablement changes are effected only after FXOS reboot.
Confirm change? (yes/no) [yes]:yes
firepower-4245 /fabric-interconnect # connect local-mgmt
firepower-4245(local-mgmt)# reboot
Before rebooting, please take a configuration backup.
Do you still want to reboot? (yes/no):yes
```

```
Broadcast message from admin@firepower-4245 (Wed Feb 21 05:59:55 2024):
All shells being terminated due to system /sbin/reboot
```

Step 2 Enable the USB port and reboot for the change to take effect.

a) Enable the USB port.

scope fabric-interconnect

enable usb-port

commit buffer

b) Reboot the chassis.

connect local-mgmt

reboot

Example:

```
firepower-4245 /fabric-interconnect # enable usb-port
Note: USB enablement or disablement changes are effected only after FXOS reboot.
Confirm change? (yes/no) [yes]:
device /fabric-interconnect* # commit buffer
Note: USB enablement or disablement changes are effected only after FXOS reboot.
Confirm change? (yes/no) [yes]:yes
firepower-4245 /fabric-interconnect # connect local-mgmt
firepower-4245(local-mgmt)# reboot
Before rebooting, please take a configuration backup.
Do you still want to reboot? (yes/no):yes
Broadcast message from admin@firepower-4245 (Wed Feb 21 05:59:55 2024):
All shells being terminated due to system /sbin/reboot
```

Step 3 View the USB port status.

scope fabric-interconnect

show usb-port

The Admin State shows the USB port configuration. The Oper State shows the current operation. For example, if you disable the USB port but do not reload, the Admin State will show Disabled while the Oper State would will Enabled.

Example:

```
firepower-4245# scope fabric-interconnect
firepower-4245 /fabric-interconnect # show usb-port
Usb Port:
Equipment         Admin State Oper State
----------------- -----------  ----------
A                 Disabled    Disabled
```

History for Device Management

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Add device by registration key using basic initial configuration added to the Device (Wizard)	7.7.0	Any	You can now use the Device (Wizard) to add a device using a registration key with a basic initial configuration. This functionality is still present on the **Add** > **Device** screen as well. New/modified screens: **Devices** > **Device Management** > **Add** > **Device (Wizard)**
Serial-number registration (zero-touch provisioning) supported from an on-prem Firewall Management Center.	7.6.0	Mgmt. center must be publicly reachable: 7.2.0 Restriction removed: 7.2.4/7.4.0	You can now register a device using its serial number from an on-prem Firewall Management Center. With templates (requires Firewall Threat Defense 7.4.1+ on the device), you can register multiple devices at once. This feature was previously known as low-touch provisioning. Requires Cisco Security Cloud. For upgraded Firewall Management Centers, your existing Security Cloud Control integration continues to work until you enable Cisco Security Cloud. New/modified screens: **Devices** > **Device Management** > **Add** > **Device (Wizard)** Supported platforms: Firepower 1000/2100, Secure Firewall 1200/3100. Note that Firepower 2100 support is for Firewall Threat Defense 7.4.1–7.4.x only; those devices cannot run Version 7.6.0.
Delete menu item renamed to **Unregister**	7.6.0	Any	The **Delete** menu choice was renamed to **Unregister** to better indicate that the device, high-availability pair, or cluster is being unregistered from the Firewall Management Center and not deleted from the high availability pair or cluster or having its configuration erased. The device, high-availability pair, or cluster continues to pass traffic until it is re-registered. New/modified screens: **Devices** > **Device Management** > **More**
Add devices using templates	7.6.0	7.4	The **Devices** > **Device Management** > **Add** > **Device (Wizard)** screen lets you add devices using a template. New/modified screens: **Devices** > **Device Management** > **Add** > **Device (Wizard)**
Disable the front panel USB-A port on the Firepower 1000 and Secure Firewall 3100/4200.	7.6.0	7.6.0	You can now disable the front panel USB-A port on the Firepower 1000 and Secure Firewall 3100/4200. By default, the port is enabled. New/modified Firewall Threat Defense CLI commands: **system support usb show**, **system support usb port disable**, **system support usb port enable** New/modified FXOS CLI commands for the Secure Firewall 3100/4200 in multi-instance mode: **show usb-port**, **disable USB port**, **enable usb-port** See: Cisco Secure Firewall Threat Defense Command Reference and Cisco Firepower 4100/9300 FXOS Command Reference

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Chassis-level health alerts for the Firepower 4100/9300.	7.4.1	7.4.1	You can now view chassis-level health alerts for Firepower 4100/9300 by registering the chassis to the Firewall Management Center as a read-only device. You must also enable the Firewall Threat Defense Platform Faults health module and apply the health policy. The alerts appear in the Message Center, the health monitor (in the left pane, under Devices, select the chassis), and in the health events view. You can also add a chassis (and view health alerts for) the Secure Firewall 3100 in multi-instance mode. For those devices, you use the Firewall Management Center to manage the chassis. But for the Firepower 4100/9300 chassis, you still must use the chassis manager or the FXOS CLI. New/modified screens: **Devices** > **Device Management** > **Add** > **Chassis**
Zero-Touch Provisioning to register the Firepower 1000/2100 and Secure Firewall 3100 to the Firewall Management Center using a serial number.	7.4.0	Mgmt. center *is* publicly reachable: 7.2.0 Mgmt. center *is not* publicly reachable: 7.2.4/7.4.0	Zero-Touch Provisioning (also called low-touch provisioning) lets you register Firepower 1000/2100 and Secure Firewall 3100 devices to the Firewall Management Center by serial number without having to perform any initial setup on the device. The Firewall Management Center integrates with SecureX and Security Cloud Control for this functionality. New/modified screens: **Devices** > **Device Management** > **Add** > **Device** > **Serial Number** Version restrictions: This feature is not supported on Version 7.3.x or 7.4.0 Firewall Threat Defense when the Firewall Management Center is not publicly reachable. Support returns in Version 7.4.1.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Merged management and diagnostic interfaces.	7.4.0	7.4.0	**Upgrade impact. Merge interfaces after upgrade.** For new devices using 7.4 and later, you cannot use the legacy diagnostic interface. Only the merged management interface is available. If you upgraded to 7.4 or later and: • You did not have any configuration for the diagnostic interface, then the interfaces will merge automatically. • You have configuration for the diagnostic interface, then you have the choice to merge the interfaces manually, or you can continue to use the separate diagnostic interface. Note that support for the diagnostic interface will be removed in a later release, so you should plan to merge the interfaces as soon as possible. Merged mode also changes the behavior of AAA traffic to use the data routing table by default. The management-only routing table can now only be used if you specify the management-only interface (including Management) in the configuration. For platform settings, this means: • You can no longer enable HTTP, ICMP, or SMTP for diagnostic. • For SNMP, you can allow hosts on management instead of diagnostic. • For Syslog servers, you can reach them on management instead of diagnostic. • If Platform Settings for syslog servers or SNMP hosts specify the diagnostic interface by name, then you must use separate Platform Settings policies for merged and non-merged devices. • DNS lookups no longer fall back to the management-only routing table if you do not specify interfaces. New/modified screens: **Devices** > **Device Management** > **Interfaces** New/modified commands: **show management-interface convergence**
Migrate Firepower 1000/2100 to Secure Firewall 3100.	7.4.0	Any	You can now easily migrate configurations from the Firepower 1000/2100 to the Secure Firewall 3100. New/modified screens: **Devices** > **Device Management** > **Migrate** Platform restrictions: Migration not supported from the Firepower 1010 or 1010E.
Download a report of all registered devices.	7.4.0	Any	You can now download a report of all registered devices. On **Devices** > **Device Management**, click the new **Download Device List Report** link, at the top right of the page.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Manage Firewall Threat Defense high availability pairs using a data interface.	7.4.0	7.4.0	Firewall Threat Defense high availability now supports using a regular data interface for communication with the Firewall Management Center. Previously, only standalone devices supported this feature. See: Device Management
ISA 3000 System LED support for shutting down.	7.0.5/7.3.0	7.0.5/7.3.0	When you shut down the ISA 3000, the System LED will turn off. You should wait at least 10 seconds before removing the power.
ISA 3000 support for shutting down.	7.0.2/7.2.0	7.0.2/7.2.0	You can now shut down the ISA 3000; previously, you could only reboot the device.
Multi-manager support.	7.2.0	7.2.0	We introduced the cloud-delivered management center. The cloud-delivered management center uses the Security Cloud Control Firewall Management (Security Cloud Control) platform and unites management across multiple Cisco security solutions. We take care of manager updates. Hardware or virtual management centers running Version 7.2+ can "co-manage" cloud-managed devices, but for event logging and analytics purposes only. You cannot deploy policy to these devices from the hardware or virtual management center. New/modified commands: **configure manager add**, **configure manager delete**, **configure manager edit**, **show managers** New/modified screens: • When you add a cloud-managed device to a hardware or virtual management center, use the new **Security Cloud Control Managed Device** check box to specify that it is analytics-only. • View which devices are analytics-only on **Devices** > **Device Management**. For more information, see Security Cloud Control documentation.
RAID support for SSDs on the Secure Firewall 3100.	7.1.0	7.1.0	The SSDs are self-encrypting drives (SEDs), and if you have 2 SSDs, they form a software RAID. New/modified commands: **configure raid, show raid, show ssd**
Support for TLS 1.3 for the management connection.	7.1.0	7.1.0	The FMC-device management connection now uses TLS 1.3. Previously, TLS 1.2 was supported.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Use FDM to configure FTD for management by the FMC.	7.1.0	7.1.0	When you perform initial setup using FDM, all interface configuration completed in FDM is retained when you switch to FMC for management, in addition to the Management and manager access settings. Note that other default configuration settings, such as the access control policy or security zones, are not retained. When you use the FMC CLI, only the Management and manager access settings are retained (for example, the default inside interface configuration is not retained). After you switch to FMC, you can no longer use FDM to manage FTD. New/modified FDM screens: **System Settings** > **Management Center**
Filter devices by upgrade status.	6.7.0	6.7.0	The **Device Management** page now provides upgrade information about your managed devices, including whether a device is upgrading (and what its upgrade path is), and whether its last upgrade succeeded or failed. New/modified screens: **Devices** > **Device Management**
One-click access to the Firepower Chassis Manager.	6.4.0	6.4.0	For Firepower 4100/9300 series devices, the Device Management page provides a link to the Firepower Chassis Manager web interface. New/modified screens: **Devices** > **Device Management**
Filter devices by health and deployment status; view version information.	6.2.3	6.2.3	The Device Management page now provides version information for managed devices, as well as the ability to filter devices by health and deployment status. New/modified screens: **Devices** > **Device Management**

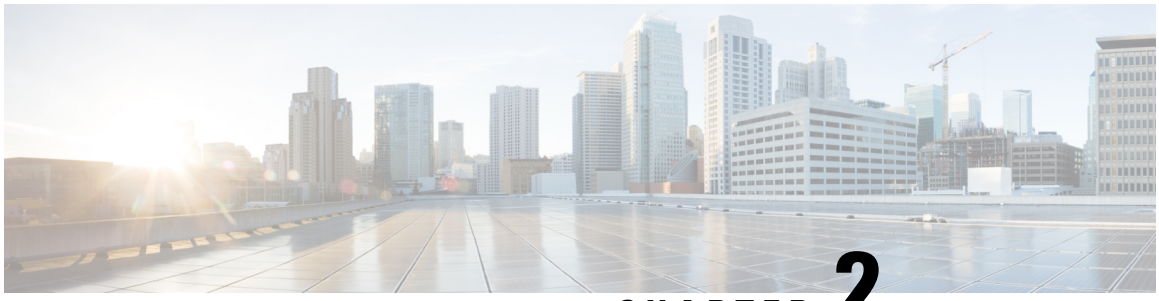

CHAPTER 2

Device Management Using Device Templates

You can add devices to the Firewall Management Center using device templates.

- About Device Management using Device Templates, on page 89
- Requirements and Prerequisites for Device Management using Device Templates, on page 93
- Licenses for Device Management using Device Templates, on page 94
- Guidelines and Limitations for Device Management using Device Templates, on page 94
- Workflow for Device Management using Device Templates, on page 98
- Configure Device Templates, on page 98
- Use Templates with Devices, on page 119
- Monitoring Device Templates, on page 120
- Troubleshooting Device Templates, on page 122
- History for Device Management using Device Templates, on page 125

About Device Management using Device Templates

Device templates enable deployment of multiple branch devices with pre-provisioned initial device configurations. You can use device templates to perform bulk zero-touch provisioning of multiple devices, apply day 2 configuration changes to multiple devices with different interface configurations, and clone configuration parameters from existing devices. You can also register more than one device at a time with the Firewall Management Center using serial numbers.

When you register a device using basic initial configuration, you can apply limited configurations such as the access control policy and licenses. You must then configure other device settings such as interfaces, routing, and site-to-site VPN configurations individually after device registration. Device templates let you pre-configure these settings and more so you can apply them at the time of registration. Values that need to be unique per device, such as IP addresses, can be defined using variables and network object overrides that you define at registration.

You can also configure site-to-site VPN connections in a device template. These configurations define the site-to-site VPN topologies that a device should be a part of. The VPN configurations along with the other device template policies and configurations enable easy deployment of the branch device to your network. Device templates support the configuration of a device only as a spoke. A device can be part of multiple hub and spoke site-to-site VPN topologies.

After the configured device template is applied to a device, the variables are resolved, the protected network overrides are configured, and the device is added as a spoke in the specified VPN topology.

Methods to Register Devices using Templates

You can use device templates with the following methods to register the device on the Firewall Management Center and set up day 0 configuration:

- Registration Key - You can register a single device by specifying the registration key and defining variables in the Firewall Management Center.
- Serial Number - You can use zero-touch provisioning to register one or more devices by serial number. For serial number registration, define all variables and overrides in a CSV file that you upload.

Variables and Network Object Overrides

You can parameterize template configurations using variables and network object overrides.

A variable is an object type that is supported for template configurations. A variable in a template defines specific configuration values for a device. You can define values for these variables during device registration and during application of the template on the device. You can see the variable icon (x) for the fields that use a variable. The variables are displayed with a $ prefix to distinguish these values from the other values.

For information on supported variable types and creating variables, see Supported Variables and Add a Variable.

Network object overrides are similar to variables. But, these are used to provide override values for a network object. You can declare a list of network objects in the template and create network object overrides for these objects. You can then provide values for these network object overrides during the application of the template on the device. For example, if you define a host network object in a template, you can add a network object override before the application of the template on the device and then provide a relevant value during the application of the template on the device.

For more information on supported network objects and adding a network object override, see Supported Network Object Overrides and Add a Network Object Override.

Model Mapping

As interface configurations vary for different device models, the interface configurations in the template have to be copied to the target interfaces on the device. Model mapping enables you to define mapping of interfaces defined in the template to the interfaces of the required Firewall Threat Defense model. During application of the template on the device, the variables in the interface configurations are replaced with the values that you provide and copied to the mapped interfaces on the device. Note that you have to create the model mappings in the template before initiating application of the template on the device. For more information on setting up model mapping, see Add Model Mapping.

Templates and High Availablity

You can apply device templates on Firewall Threat Defense high-availability devices after device registration. High-availability-specific configurations are not supported in device templates. Any high-availability configurations and monitored interfaces that are already part of the target high-availability device pair configurations are not modified. You cannot map any template interfaces to failover interfaces.

You can generate a device template from a high-availability device pair. Template operations such as application of the template on a device, template generation, importing, and exporting of a template, can be performed only on the active unit. You cannot perform these operations on the standby unit.

Templates and Domains

Device templates can exist in any domain. If you are in the child domain, you have read-only access to the templates above you in the domain hierarchy. You can apply a template to a device from its domain or its parent domains. You can generate a template from a device and apply that template to a device in any domain in the domain hierarchy.

A domain hierarchy sample is given below along with a table displaying the supported device template application and generation scenarios.

Consider the following scenario:

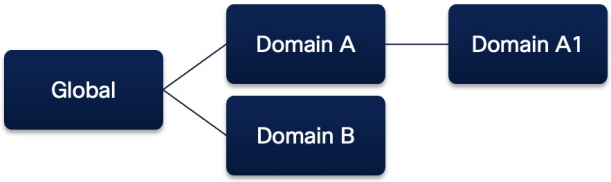

- Domain A and B are child domains of the Global domain.
- Domain A1 is the child domain of Domain A.

Template Domain	Device Domain	Device Template Application/Generation Supported
Global	A1	Yes
Global	B	Yes
A	A1	Yes
A	B	No
B	A1	No
B	B	Yes
A1	A1	Yes
A1	B	No

Domains and VPN Connections

- You can define a template in a global or child/leaf domain. However, you can define a VPN topology only in a leaf domain.
- You can configure VPN connections in a template for all domains. During template application, VPN connections are applied to the device only if the device is in the same domain as the VPN topology.

A domain hierarchy sample is given below along with a table displaying the supported device template application and generation scenarios.

Consider the following scenario:

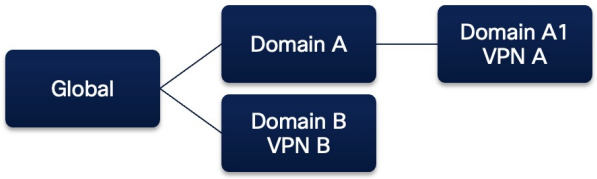

- Domain A and B are child domains of the Global domain.
- Domain A1 is the child domain of Domain A.
- VPN A is part of Domain A1.
- VPN B is part of Domain B.

Template Domain	VPN Topology in the Template	Device Domain	Device Template Application/Generation Supported
Global	VPN A VPN B	A1	No
Global	VPN B	B	Yes
A	VPN A	A1	Yes
B	VPN B	A1	No
B	VPN B	B	Yes
A1	VPN A	A1	Yes

Validation of Template Configuration Before and After Application of Template on Device

Validation of template configuration is done before and after application of the template on the device.

The following validation checks are performed at the start of the task to apply the template on the device:

- Ensure that the target device model and version are supported.
- Cluster and container checks -The device must not be part of a cluster or multi-instance.
- Model mapping validation - Model mapping for the target device model exists and is valid.
- Sanity check of template parameter values. For example, two variables used as IP addresses of interfaces must not have the same value.

The following validation checks are performed at the end of the task to apply the template on the device to ensure that the applied configurations are valid:

- Interface configuration validation. For example, variables used for the IP address fields of two or more interfaces must not have the same IP address values.
- Routing policy validation. For example, the IPv4 address in BGP neighbor configurations must not overlap with the IP address of any interface.

If the validation checks that are done at the end of the task to apply the template on the device fail, any applied configurations are rolled back and the device is restored to it's original state.

Requirements and Prerequisites for Device Management using Device Templates

Model Support

Device templates are supported on On-Prem Firewall Management Center, Cloud-Delivered Firewall Management Center(cdFMC), with the following models running Secure Firewall version 7.4.1 and later versions:

- Firepower 1010
- Firepower 1100
- Secure Firewall 1200
- Firepower 2100 (7.4.x only)
- Secure Firewall 3100

Supported Domains

Any

User Roles

- To create, modify, or delete templates:
  - Admin
  - Network Admin
- To view the created templates:
  - Any

Prerequisites for VPN Connections in Device Templates

- Configure site-to-site VPN topologies that must be used in the device template.

- Ensure that you have configured all hub and VPN topology-related configurations such as authentication methods, IKE and IPsec policies.
- Supported types of VPN hub and spoke topologies are:
  - Policy-based
  - Route-based
  - SD-WAN
- Assign appropriate logical names and IP addresses to the interfaces of the Firewall Threat Defense devices. For example, use *inside* for the interface connected to the LAN, and *outside* for the interface connected to the internet or WAN.
- Spoke devices must be version 7.4.1 and later.

Licenses for Device Management using Device Templates

- Device templates does not have any specific license requirements.
- License entitlements for the target device must be present in the Smart Licensing account.
- To configure VPN connections in the template, the Essentials license must allow export-controlled functionality. Choose **System** (⚙) > **Licenses** > **Smart Licenses** to verify this functionality in the Firewall Management Center.
- When you apply a template on a device, note the following conditions for Secure Client licensing:

Device with Secure Client License	Template with Secure Client License	Secure Client License after Device Template Application
Yes	Yes	Template License
Yes	No	Device License
No	Yes	Template License

Guidelines and Limitations for Device Management using Device Templates

General Guidelines for Device Templates

- All device configurations other than VNI and VTEP are supported.
- You can attach shared policies and S2S VPN policies to a template. These policies are assigned during template application.
- Templates can be applied on HA devices. However, application of device templates during HA device pair registration is not supported. You also cannot manage HA-related configurations such as failover

links, standby IP addresses, and so on. For more information, see Device Template Operations on Threat Defense HA Devices.

- Device template operations are supported only on the active Firewall Management Center. The standby peer does not support device template operations.
- Ensure that the template names and device display names are not the same.
- Ensure that you do not create or delete a template during device backup or restore operations.
- After application of the template on the device, if the manager access is changed from management to data interface or vice-versa, you must re-establish the management connection with the device. Note that you cannot change the manager access interface during template application.
- You can add a maximum of 250 device templates to the Firewall Management Center.
- Templates that are created and configured for devices that are managed through the data interface cannot be used to register and apply to devices that are managed through the management interface.
- Device registration and application of template does not come under the Change Management workflow. Only approved data, such as access policies, templates, template variables, network overrides declared in template, and template configurations used in the template application operation are used.
- Device registration with serial number and access control policy is supported for only one device at a time.
- Devices with IPv6 DHCP discoverability are not supported when you add devices using serial numbers.
- When you apply a template to a device that is already registered, the configuration in the template is only copied to the target device. You can then choose to manually deploy the configuration on the device or let the copied configuration stay on the device and deploy later. However, if you apply the template during device onboarding, the configuration in the template is copied to the target device and, as per existing behavior, automatically deployed on the device after device registration.
- Any change in model mapping causes the device to be marked as ‘Out of Sync’. Consider reapplying the template to the devices if you have made changes to the interface mapping in the corresponding model mapping or if you have made any configuration changes after the previous application of the template.
- Device templates are only supported on merged management and diagnostic interfaces. For more information, see Merge the Management and Diagnostic Interfaces.
- Template configuration updates are supported by change management. Creation and application of device templates is not supported by change management.
- You cannot sync configuration changes from a device to the template. A few sample scenarios during which you may want to make changes to the device configuration using a template along with recommended solutions are given below:
  - If you want to test the new configuration changes on one device before propagating the changes to multiple devices, we recommend that you make the changes in the template and apply the template to one device. Validate the changes on that device and then apply the template to the other devices.
  - If you want to make a large number of changes resulting in a signification deviation from the current configuration on a device, and then propagate those changes to other devices, you may choose one of the following options:

- Export the current template to get a copy of the template. You can then make the required changes in the template and apply to a single device. Validate the changes on that device and then apply the template to the other devices.
- You can also make the required changes on the device and create a template from that device. You can then apply and validate the changes on the other devices. However, we do not recommend this as template parameters such as variables and network object overrides will not be present in the created template.

- If the configuration on a particular device starts to differ significantly from the configuration in the template, you may also choose to not use the template for this device and delete the device-template association from the **Associated Devices** window.

Guidelines for VPN Connections in Device Templates

- Supported interfaces for VPN topologies are:

Topology Type	Interface Type

Policy-Based and SD-WAN	• Physical interfaces
	  • Non-management
	  • Interface Mode must be either Routed or None
	• Subinterfaces
	• Redundant interfaces
	• Etherchannel interfaces
	• VLAN interfaces

Route-Based	Static Virtual Tunnel Interfaces

- When you apply a template on a device that is part of a VPN topology, you must ensure that the template includes interface configurations for all interfaces used in the topology.
- When you apply a template with VPN connections to multiple devices, note the following:

  A template is applied to multiple devices in the order in which you have selected the devices. If the template has VPN connections, the corresponding VPN topology is locked.
- For SD-WAN topology VPN connections: Ensure that IP address subnet of the interface does not conflict with the subnet of the IP address pool of the SD-WAN hub.
- Domain:
  - You can define a template in a global or leaf domain. However, you can define a VPN topology only in a leaf domain.
  - You can configure VPN connections in a template for all domains. During template application, VPN connections are applied to the device only if the device is in the same domain as the VPN topology.

For more information, see Templates and Domains, on page 91.

- Change management: Before you apply a device template to a device, ensure that the VPN topology is not locked by a Change management ticket.

Limitations for Device Templates

- The following features and configurations are not supported using device templates:
  - Multi-instance mode
  - Clustering
  - Non-converged management interface
  - Transparent mode
  - HA failover configurations
  - Chassis configurations
  - Logical devices
  - Variables for nested objects
  - Override support for network groups and other object types

Limitations for VPN Connections

- When you create a template from a device that is part of a VPN topology (**Devices** > **Device Management** **More** (⋮) > **Generate Template from Device**), VPN configurations are not part of the template. You must reconfigure the VPN configurations on the template.
- When you export a device template with one or more VPN connections (**Template Settings** > **General** > **General** pane > **Export**), the VPN connections are not exported. You must reconfigure the VPN connections on the imported template.
- Certificate-based authentication:
  - Device templates do not support automatic certificate enrolment of a device.
  - When you onboard a device using a template with VPN configurations, if the VPN topology uses certificate-based authentication, the first deployment to the device will fail. Ensure that you manually enroll the device certificate after the device registration and deploy the configurations on the device again.

Workflow for Device Management using Device Templates

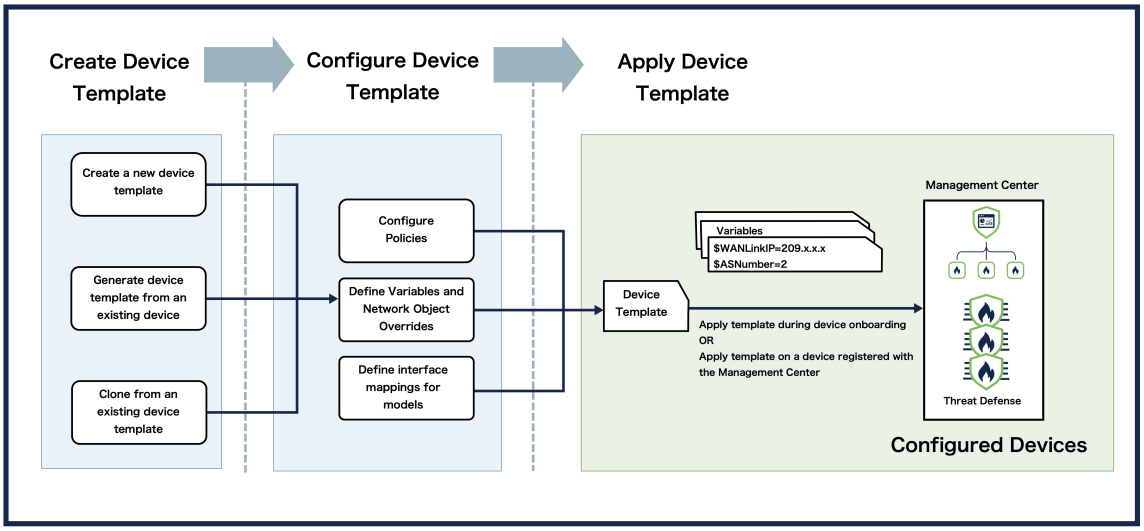

Configure Device Templates

Add a template and then configure settings in the template.

Add a Device Template

You can add a new device template with the required configuration or generate a device template from an existing device.

Create a New Device Template

You can specify the device template name, description, access control policy, and routing mode. You can add more configurations after the template has been created. Perform the procedure given below to create a device template.

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click **Add Device Template**.

Step 3 In the **Add Device Template** window, enter a **Name** for the template.

Step 4 (Optional) Enter a **Description** for the template.

Step 5 Choose an **Access Control Policy** from the drop-down list.

Step 6 Choose a **Mode** from the drop-down list.

Step 7 Click **OK**.

Generate a New Device Template from an Existing Device

You can generate a new device template from a device that is registered with the Firewall Management Center. The new template has the same configuration as the device from which it is generated. You can generate a new device template from standalone and HA devices. However, if you generate a template from HA devices, the new template will not contain the failover configurations.

Perform the procedure given below to generate a new device template from an existing device.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click the **More** (⋮) icon, and click **Generate Template from Device**.

Step 3 In the **Generate template from device** window, enter a **Name** for the template.

Step 4 (Optional) Enter a **Description** for the template.

Step 5 Choose an **Access Control Policy** from the drop-down list.

Note
This policy is assigned to the generated template. Any other shared policies that are associated with the device from which the template is generated are assigned to the generated template only if these policies are visible in the domain in which the template is being generated.

Step 6 Click **OK**. You can view the status of template creation in the **Notifications > Tasks** window.

Step 7 Choose **Devices** > **Template Management** to view the newly created template.

Import a Device Template

You can import a template into the Firewall Management Center or export a template to your local system. This feature is useful in the following scenarios:

- Generate a copy of the template from a device, export it, and import that template into another Firewall Management Center or Cloud-Delivered Firewall Management Center.
- Generate a copy of the template, export it, modify as required to create a variation of the existing template, and import the template into the Firewall Management Center.
- Generate a copy of the template, export it, and import that template into another domain in which the source template is not visible.

When you import a template into a domain, any objects that are part of the configuration are either newly created or reused if the objects with the same name are visible in the domain into which the template is imported. Any object with matching names that is not visible, due to domain hierarchy, is imported as a new object with the name suffixed with a _x.

If there is a mismatch in the variable names when you want to onboard devices in a domain using the cloned template from another domain, you must specify the new variable names in the .csv file to onboard the devices.

Perform the procedure given below to import a device template into the Firewall Management Center from your local system.

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the **More** (⋮) icon for the template that you want to replace with an imported template.

Step 3 Click **Import** and click **Import** again.

If you want to export a template, click **Export** and click **OK**.

Step 4 View the status of the import task in the **Notifications > Tasks** window.

Step 5 Choose the template SFO file on your local system and click **Open**. This template SFO file that you import can be newly created, generated from a device, or cloned from an existing template.

Step 6 View the status of the import task in the **Notifications > Tasks** window. You will see a notification informing you that the import or export task is successfully completed. If you are exporting a template, click **Download Export Package** in the **Notifications > Tasks** window to download the template configuration as an SFO file.

Note
Alternatively, you can also go to the **Template Management** window and click the **Edit** (✎) icon of the template. Then, go to **Template Settings > General**, and click **Import** or **Export** in the **General** tile to import or export the template.

Configure Device Settings in the Template

After creating the template, you can set up device configurations and configure settings that you want to apply on the device by editing the template.

Add a Physical Interface

By default, a device template will enable the device to come up with the following physical interfaces:

- Management interface
- Inside interface
- Outside interface

Perform the procedure given below to create a physical interface.

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the **Edit** (✎) icon of the template in which you want to add the physical interface.

Step 3 In the **Interfaces** tab, click **Add Physical Interface**.

Step 4 Choose a **Slot** and **Port Index** number from the drop-down list.

Step 5 Click **Create Interface**.

Add a Logical Interface

You can create a logical interface in the same way as you do on the Firewall Management Center without using the template. Perform the procedure given below to create a logical interface.

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the **Edit** (✎) icon of the template in which you want to add the logical interface.

Step 3 In the **Interfaces** tab, click **Add Interface**, and choose the type of interface that you want to create from the drop-down list. You can create the following types of interfaces:

- Sub-interface
- Ether channel interface
- Bridge group interface
- VLAN interface
- Virtual tunnel interface
- Loopback interface

For more information, see Interface overview and Regular Firewall Interfaces.

Edit an Interface

You can edit an interface in the same way as you do on the Firewall Management Center without using the template. Use template variables to set up the IPv4 and IPv6 addresses. The device template supports the configurations that are supported on Firepower 1000,Secure Firewall 1200, Firepower 2100, and Secure Firewall 3100 Firewall Threat Defense devices. Perform the procedure given below to edit an interface.

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the **Edit** (✎) icon of the template in which you want to edit the physical interface.

Step 3 In the **Interfaces** tab, click the **Edit** icon for the interface that you want to edit.

Step 4 In the **Edit Physical Interface** window, you can edit any of the following settings:

- General

- PoE
- IPv4
- IPv6
- Path Monitoring
- Hardware Configuration
- Manager Access
- Advanced

Note
Use variables to configure IPv4 and IPv6 addresses. For more information on templatizing variables, see Configure Template Parameters.

For more information on editing the settings mentioned above, see Interface overview and Regular Firewall Interfaces.

Configure Other Device Settings

Configure the other device settings in the same way as you do on the Firewall Management Center without using the template.

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the **Edit** (✎) icon for the template in which you want to configure the settings.

Step 3 Click the tabs at the top of the window to configure any of the following settings:

- Inline Sets
- Routing
- DHCP
- VPN
- Template Settings

Configure Template Settings

These are template-specific settings that are copied to the device when the template is applied on the device. In the **Template Settings** window, you can configure the following template settings:

- General

- General
- License
- Applied Policies
- Advanced Settings
- Deployment Settings

- Template Parameters
  - Variables
  - Network object overrides

- Model Mapping

Edit General Settings

In the **General** tile, you see the following fields:

- **Template Name** – The name of the template.
- **Transfer Packets** – Displays whether or not the managed device sends packet data with the events to the management center.
- **Mode** – Displays the mode of the management interface for the device: routed.
- **Configuration** – Click **Export** to export the template configurations as an SFO file. Click **Import** to import an SFO file that has the template configurations that you require.
- **Manage device by Data Interface** – Toggle the button to enable or disable management of the device using the data interface.

Perform the procedure given below to edit the name of the device, and to enable or disable packet transfer.

Procedure

Step 1 Click the **Edit** (✎) icon in the **General** tile.

Step 2 Change the **Template Name** as per your requirement.

Step 3 Check the **Transfer Packets** checkbox to allow packet data to be stored with events on the Firewall Management Center.

Step 4 Click **Save**.

Edit Licenses

In the **License** tile, you can see the **License types** that are required based on the configurations used in the template. Choosing a license here does not consume that license on the device. The license is consumed only when you apply the template to a device.

Perform the procedure given below to edit the license types as per your requirement.

Procedure

Step 1 Click the **Edit** (✎) icon in the **License** tile.
Step 2 Check or clear the check box next to the license you want to enable or disable for the managed device.
Step 3 Click **Save**.

Edit Applied Policies

In the **Applied Policies** tile, you can see the access control policies that are associated with the template.

For policies with links, you can click the link to view the policy.

Perform the procedure given below to edit the policy assignments as per your requirement.

Procedure

Step 1 Click the **Edit** (✎) icon in the **Applied Policies** tile.
Step 2 For each policy type, choose a policy from the drop-down list. Only existing policies are listed.
Step 3 Click **Save**.

Edit Advanced Settings

The **Advanced Settings** tile displays the advanced configuration settings, as described below. You can edit any of these settings.

Table 2: Advanced Section Table Fields

Field	Description
Application Bypass	The state of Automatic Application Bypass on the device.
Bypass Threshold	The Automatic Application Bypass threshold, in milliseconds.

Field	Description
Object Group Search	The state of object group search on the device. While operating, the Firewall Threat Defense device expands access control rules into multiple access control list entries based on the contents of any network or interface objects used in the access rule. You can reduce the memory required to search access control rules by enabling object group search. With object group search enabled, the system does not expand network or interface objects, but instead searches access rules for matches based on those group definitions. Object group search does not impact how your access rules are defined or how they appear in the Firewall Management Center. It impacts only how the device interprets and processes them while matching connections to access control rules. **Note** By default, the **Object Group Search** is enabled when you add Firewall Threat Defense for the first time in the Firewall Management Center.
Interface Object Optimization	The state of interface object optimization on the device. During deployment, interface groups and security zones used in the access control and prefilter policies generate separate rules for each source/destination interface pair. If you enable interface object optimization, the system will instead deploy a single rule per access control/prefilter rule, which can simplify the device configuration and improve deployment performance. If you select this option, also select the **Object Group Search** option to reduce memory usage on the device.

Perform the procedure given below to edit the advanced settings.

Procedure

Step 1 Click the **Edit** (✎) icon in the **Advanced Settings** tile.

Step 2 You can change the settings as per your requirement. For more information, see the following sections:

- Configure Automatic Application Bypass
- Configure Object Group Search
- Configure Interface Object Optimization

Step 3 Click **Save**.

Edit Deployment Settings

The **Deployment Settings** tile displays the information described in the table below.

Table 3: Deployment Settings

Field	Description
Auto Rollback Deployment if Connectivity Fails	Enabled or Disabled. You can enable auto rollback if the management connection fails as a result of the deployment; specifically if you use data for Firewall Management Center access, and then you misconfigure the data interface.
Connectivity Monitor Interval (in Minutes)	Shows the amount of time to wait before rolling back the configuration.

Deployment settings include enabling auto rollback of the deployment if the management connection fails as a result of the deployment; specifically if you use data for Firewall Management Center access, and then you misconfigure the data interface. You can alternatively manually roll back the configuration using the **configure policy rollback** command.

Perform the procedure given below to edit the deployment settings.

Procedure

Step 1 Click the **Edit** icon in the **Deployment Settings** tile.

Step 2 Set the **Connectivity Monitor Interval (in Minutes)** to set the amount of time to wait before rolling back the configuration. The default is 20 minutes.

Step 3 If a rollback occurs, see the following for next steps.

- If the auto rollback was successful, you see a success message instructing you to do a full deployment.
- You can also click **Deploy** and then go to the **Advanced Deploy** screen and click the **Preview** icon to view the parts of the configuration that were rolled back (see Deploy Configuration Changes). Click **Show Rollback Changes** to view the changes, and **Hide Rollback Changes** to hide the changes.
- In the Deployment History Preview, you can view the rollback changes.

Step 4 Check that the management connection was reestablished.

In Firewall Management Center, check the management connection status on the Connection Status page. Navigate to **Devices** > **Device Management** page and then navigate to **Management** area under the **Device** tab. Then in the **FMC Access Details** screen, click **Connection Status**.

At the Firewall Threat Defense CLI, enter the **sftunnel-status-brief** command to view the management connection status.

If it takes more than 10 minutes to reestablish the connection, you should troubleshoot the connection. See Troubleshoot Management Connectivity on a Data Interface.

Configure Template Parameters

You can templatize configurations using template parameters such as variables and network object overrides.

Supported Variables

The following variable types are supported in device templates.

Variable Name	Description	Type
AS Number	Defines the unique Autonomous System (AS) number.	Integer Example: 2
FQDN	Defines a single Fully Qualified Domain Name (FQDN).	String Example: abc.example.com
IPv4 Host	Defines the IPv4 address of the host.	String Example: 209.165.201.8
IPv4 Network	Defines the IPv4 network address block.	String Example: 209.165.200.224/27
IPv4 Range	Defines the range of IPv4 addresses.	String Example: 209.165.200.225-209.165.200.250
IPv6 Host	Defines the IPv6 address of the host.	String Example: 2001:DB8::1
IPv6 Network	Defines the IPv6 network address block.	String Example: 2001:DB8:0:CD30::/60
Password	Defines a password string.	String Example: E28@2OiUrhx!
Router ID	Defines an identifier for the router.	Integer Example: 21
String	Defines a custom string.	String Example: testvalue2

Add a Variable

Perform the procedure given below to add a variable.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Variable** from the list of object types.

Step 3 Click **Add Variable**.

Step 4 Enter a **Name**.

In a multi-domain deployment, object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

Step 5 Choose a **Variable Type** from the drop-down list.

Step 6 (Optional) Enter a **Description**.

Step 7 Click **Save**.

Supported Network Object Overrides

The following network objects are supported.

Network Object Name	Description	Type
Network	An address block, also known as a subnet.	String Example: IPv4 - 209.165.200.224/27 IPv6 - 2001:DB8::/48
Host	The IP address of the host.	String Example: IPv4 - 209.165.200.225 IPv6 - 2001:DB8:1::1
Range	A range of IP addresses.	String Example: IPv4 - 209.165.200.225-209.165.200.250 IPv6 - 2001:DB8::1 - 2001:DB8:FFFF:FFFF:FFFF:FFFE:FFFF:FFFF
FQDN	A single fully-qualified domain name (FQDN).	String Example: abc.example.com

Add a Network Object Override

Perform the procedure given below to add a network object override.

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the **Edit** (✎) icon of the template in which you want to add the network object override.

Step 3 Choose **Template Settings** > **Template Parameters**.

Step 4 In the **Network Object Overrides** section, click **Add or Remove Network Object Overrides**.

Step 5 In the **Add or Remove Network Object Overrides** window, choose the network objects for which you want to create network object overrides from the **Available Networks** window and click the > button.

Step 6 Click **Save**.

Add Model Mapping

For each model, you can specify which template interface corresponds with which model interface. You can map a template to one or more models as long as the interface configurations are valid for all mapped models. For example, if the template includes switch ports and VLAN interfaces, then that template can only be applied to a Firepower 1010 or Secure Firewall 1210/1220.

Perform the procedure given below to add model mapping.

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click **Add Model Mapping** for the template in which you want to create the model mapping. Alternatively, you can click the **Edit** (✎) icon of the template and choose **Template Settings > Model Mapping**.

Step 3 Click **Add Model Mapping**.

Step 4 Choose the **Device Model** from the drop-down list.

Step 5 Map the template interfaces to the device model interfaces by choosing the interface from the **Model Interface** drop-down list.

Note
You can click **Clear Mapping** to remove the defined model mapping. Click **Reset Mappings** for default interface mapping in which the mapping is done based on the slot and port index order of the interface names.

Step 6 Click **Save**. The interface mappings are listed along with the device model and mapping status on the **Model Mapping** window.

Note
Some configurations in the template may not be supported on all device models. Unsupported configurations, if any, are not applied to the device. The Device Template Apply Report provides details about such configurations.

Invalid Model Mappings

Some configurations in the template may not be supported on all device models. Unsupported configurations, if any, are not applied to the device. Valid model mappings can also become invalid when you modify template configurations. For example, when you add a new interface on the template and assign a name to it, the new interface must be mapped to the appropriate interface on the device model.

Model mapping can also be invalidated due to any of the following reasons:

- Number of configured VRF instances exceeds the limit for a specific model.

- Interfaces mapped to incompatible models, versions, or interfaces. See Requirements and Prerequisites for more information.
- Number of interfaces exceeds the model limit.
- Deleted an interface that was mapped.
- Newly added physical interfaces are not mapped to a compatible model interface.
- Model mapping is not done for a named interface.
- Model mapping is not done for an interface related to other logical interfaces, such as sub-interfaces, PC interfaces, and so on.
- Making policy or configuration changes that are unsupported on some device models. For example, enabling switch port configuration on interfaces.

You can also save a template with invalid model mappings. However, you must review and fix the model mapping before initiating application of the template on the device.

You can hover over **Invalid** under **Mapping Status** to view the errors that caused the invalid mapping status. Fix the errors before initiating application of the template on the device.

Configure Site-to-Site VPN Connections in a Device Template

Device Template with Site-to-Site VPN Connections Workflow

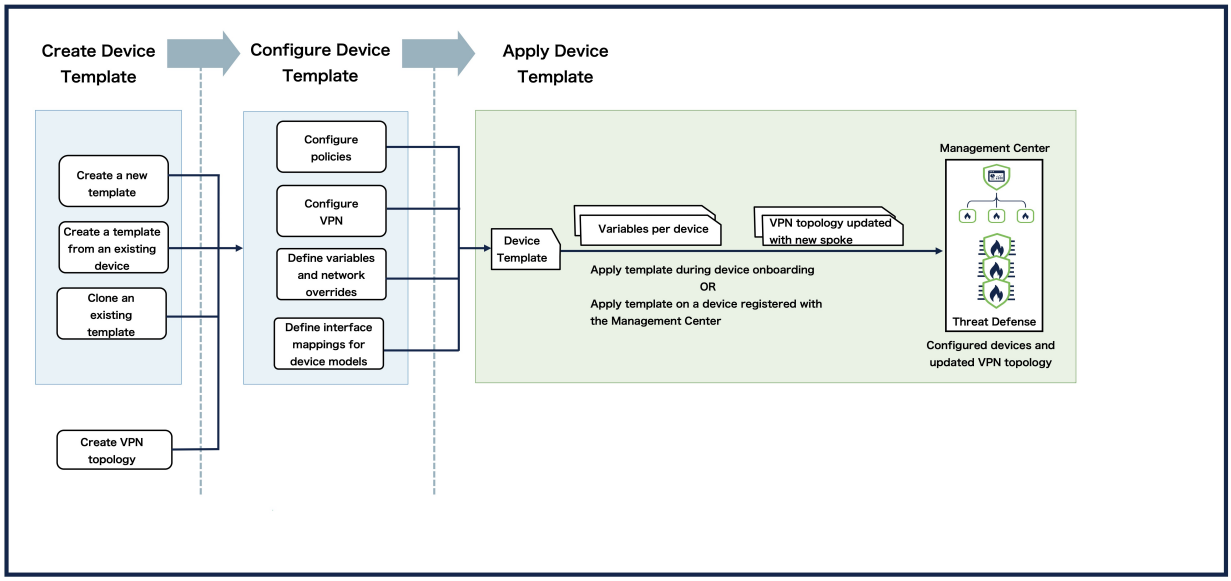

Configure an SD-WAN VPN Connection

You can configure an SD-WAN VPN connection to add spokes to SD-WAN topologies using the device template.

Before you begin

- Configure a minimum of one SD-WAN topology (**Devices** > **VPN** > **Site to Site**).
- Review Prerequisites for Configuring Device Templates and Guidelines and Limitations for Device Templates.

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the edit icon adjacent to the device template that you want to edit.

Step 3 Click the **VPN** tab.

Step 4 Click **Add VPN Connection**.

Step 5 Choose an SD-WAN topology from the **VPN Topology** drop-down list.

The **Add VPN Connection** dialog box expands and you can configure the following parameters:

a) From the **VPN Interface** drop-down list, choose a WAN-facing or internet-facing physical interface to establish a VPN connection with the hub.

This list contains all the interfaces configured on the device template.

b) **Use IP Address from the VPN Interface**—This drop-down list is auto populated with the IP address variable. For IPv6 addresses, choose an IPv6 address from the drop-down list.

c) Check the **Local Tunnel (IKE) Identity** check box to enable a unique and configurable identity for the VPN tunnel from the spoke to a remote peer.

d) **Identity Type**—Key ID is the only supported identity type. Choose a key ID variable from the drop-down list or click (+) to create a new key ID variable.

e) Click **OK**.

You can view the VPN connection in the **Site-to-Site VPN Connections** table.

Step 6 Click **Save**.

What to do next

1. Configure the routing policy for the spoke in the device template.
2. Map the device interfaces to the template interfaces (Model Mapping).
3. Apply the template to a device.

Configure a Route-Based Site-to-Site VPN Connection

You can configure a route-based site-to-site VPN connection to add spokes to route-based site-to-site VPN topologies using the device template.

Before you begin

- Configure a minimum of one route-based site-to-site VPN topology (**Devices** > **VPN** > **Site to Site**).

- Review Prerequisites for Configuring Device Templates and Guidelines and Limitations for Device Templates.

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the edit icon adjacent to the device template that you want to edit.

Step 3 Click the **VPN** tab.

Step 4 Click **Add VPN Connection**.

Step 5 Choose a route-based site-to-site VPN topology from the **VPN Topology** drop-down list.

The **Add VPN Connection** dialog box expands and you can configure the following parameters:

a) From the **Virtual Tunnel Interface (VTI)** drop-down list, choose a VTI interface or click (+) to create a new VTI.

VTI is a virtual interface used to establish a route-based VPN tunnel. You must configure routing policies for a VTI to set up a VPN tunnel. This list contains all the VTIs configured on the device template. For more information on creating a VTI, see Add a VTI Interface, on page 1537.

b) Check the **Use Public IP Address** check box to override the tunnel source IP address and configure a public IP address variable for the VTI. Click (+) to create a new public IP address variable.

This IP address is the source IP address for the VPN tunnel. By default, this is the IP address of the VPN interface. However, if the device is behind NAT, the VPN interface has a private address, but the post-NAT public IP address should be configured.

c) Check the **Local Tunnel (IKE) Identity** check box to enable a unique and configurable identity for the VPN tunnel from the spoke to a remote peer.

d) **Identity Type**: Key ID is the only supported identity type. Choose a key ID variable from the drop-down list or click (+) to create a new key ID variable.

e) (Optional) Check the **Enable Secondary VPN Tunnel** check box to configure the parameters for the secondary VPN tunnel.

f) Click **OK**.

You can view the VPN connection in the **Site-to-Site VPN Connections** table.

Step 6 Click **Save**.

What to do next

1. Configure the routing policy for the spoke in the device template.
2. Map the device interfaces to the template interfaces (Model Mapping).
3. Apply the template to a device.

Configure a Policy-Based Site-to-Site VPN Connection

You can configure a policy-based site-to-site VPN connection to add spokes to policy-based site-to-site VPN topologies using the device template.

Before you begin

- Configure a minimum of one policy-based site-to-site VPN (**Devices** > **VPN** > **Site to Site**).
- Review Prerequisites for Configuring Device Templates and Guidelines and Limitations for Device Templates.

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the edit icon adjacent to the device template that you want to edit.

Step 3 Click the **VPN** tab.

Step 4 Click **Add VPN Connection**.

Step 5 Choose a policy-based site-to-site VPN topology from the **VPN Topology** drop-down list.

The **Add VPN Connection** dialog box expands and you can configure the following parameters:

a) From the **VPN Interface** drop-down list, choose a WAN-facing or internet-facing physical interface to establish a VPN connection with the hub.

This list contains all the interfaces configured on the device template.

Do one of the following to configure the IP address of the VPN interface:

- Click the **Use IP Address from the VPN Interface** radio button to use the IP address of the VPN interface.

  This IP address is auto populated. For IPv6 addresses, choose an IPv6 address from the drop-down list.

- Click the **Use Public IP Address** radio button to configure a public IP address for the VPN interface.

  Choose an IP address variable from the drop-down list or click (+) to add an IP address variable.

b) Check the **Local Tunnel (IKE) Identity** check box to enable a unique and configurable identity for the VPN tunnel from the spoke to a remote peer.

c) **Identity Type**: Key ID is the only supported identity type. Choose a key ID variable from the drop-down list or click (+) to add a new key ID variable.

d) **Protected Networks**: Click (+) to configure a protected network for the VPN connection.

Do one of the following:

- Choose a protected network and click **OK**.
- Click **Add** to configure a network object and click **Save**.

  When you create a protected network object, note the following:

- Click either the **Host** or the **Network** radio button.
- Check the **Allow Overrides** check box.

e) Click **OK**.

You can view the VPN connection in the **Site-to-Site VPN Connections** table.

Step 6 Click **Save**.

What to do next

1. Note that before you apply a template to a device, to configure device-specific values for the protected networks, add these objects in **Template Settings > Template Parameters > Add Network Objects Overrides**.
2. Map the device interfaces to the template interfaces (Model Mapping).
3. Apply the template to a device.

Register a Device Using Device Template and Add it to a Route-Based VPN Topology

This section provides instructions to register a device using device template and add it to a route-based VPN topology.

Before you begin

Ensure that you have a route-based VPN topology.

Step	Task	GUI Path	More Information
1	Create a device template.	**Devices** > **Template Management**, and click **Add Device Template**	Create a New Device Template, on page 98
2	Configure interfaces in the template.	**Devices** > **Template Management**, and click **Interfaces**	Edit an Interface, on page 101
3	Configure a route-based VPN connection in the template.	**Devices** > **Template Management**, and click **VPN** > **Add VPN Connection**	Configure a Route-Based Site-to-Site VPN Connection, on page 111

Step	Task	GUI Path	More Information
4	Configure routing policies in the template.	**Devices** > **Template Management**, and click **Routing**	—
5	Add a model mapping for the device model in the template.	**Devices** > **Template Management**, and click **Template Settings** > **Model Mapping**	Add Model Mapping, on page 109
6	Register the device using a device template.	**Devices** > **Device Management**, and click **Add** > **Device (Wizard)**	Add a Device Using a Registration Key—Device Template, on page 36
7	Deploy configurations on the hub of the VPN topology.	**Deploy**	—

Add a Device to an SD-WAN Topology in a Dual ISP Deployment

This section provides instructions to add a device to an SD-WAN topology in a dual ISP deployment using a device template.

Before you begin

Ensure that you have two SD-WAN VPN topologies with the same hub. For more information about configuring SD-WAN topologies, see Configure an SD-WAN Topology Using the SD-WAN Wizard, on page 2016.

Step	Task	GUI Path	More Information
1	Create a device template.	**Devices** > **Template Management**	Create a New Device Template, on page 98
2	Add a physical interface in the template. By default, a template has only one outside interface. Rename the outside interfaces, for example, ISP1, ISP2.	**Devices** > **Template Management**, and click **Interfaces** > **Add Physical Interface**	Add a Physical Interface, on page 100

Step	Task	GUI Path	More Information
3	Configure an SD-WAN VPN connection using ISP1 interface.	**Devices** > **Template Management**, and click **VPN** > **Add VPN Connection**	Configure an SD-WAN VPN Connection, on page 110
4	Configure an SD-WAN VPN connection using ISP2 interface.		
5	Add static routes from ISP1 and ISP2 interfaces to the SD-WAN hub network.	**Devices** > **Template Management**, and click **Routing** > **Static Route**	-
6	Add the ISP1 and ISP2 interfaces to an ECMP zone.	**Devices** > **Template Management**, and click **Routing** > **ECMP**	-
7	Configure the network object overrides.	**Devices** > **Template Management**, and click **Template Settings** > **Template Parameters** > **Add Network Objects Overrides**	Add a Network Object Override, on page 108
8	Map the template interfaces to the device model interfaces (Model Mapping).	**Devices** > **Template Management**, and click **Template Settings** > **Model Mapping**	Add Model Mapping, on page 109
9	Apply template to the device.	**Devices** > **Template Management**	Apply a Template, on page 119
10	Deploy configurations on the device.	**Deploy**	—
11	Deploy configurations on the hubs of the SD-WAN topologies.	**Deploy**	—

For more information about dual ISP deployment using the SD-WAN wizard, see Sample Configurations for Dual ISP Deployment Using SD-WAN Wizard, on page 2020.

Configure a Template for Firewall Threat Defense Devices Managed Using the Data Interface

To configure a template that you want to apply to a Firewall Threat Defense device that is managed using a data interface for Firewall Management Center connectivity, ensure that the connectivity parameters of the device match the template. This ensures that the Threat Defense device does not lose connectivity with the Firewall Management Center after application of the template. A template that you configure for Firewall Threat Defense devices managed using the data interface cannot be applied on devices that are not managed by the data interface.

The following is a list of connectivity parameters:

- Data interface used to the manage the Firewall Threat Defense device. For example, **Ethernet1/1**.
- Name of the interface. For example, **outside**.
- IP address configured on the data interface. For example, DHCP or static IP.
- Route configured for the data interface. This can be a default or specific route defined on the data interface used for connectivity between the Firewall Threat Defense device and the Firewall Management Center.
- DDNS hostname configuration on the data interface.

If the connectivity parameters on the template do not match with the ones on the device, the template validation checks that are done to ensure that the template is successfully applied on the device will fail. The template is then not applied on the device. The template validation checks do not enforce an exact match for some parameters such as IP address or DDNS hostname. However, ensure that you configure such parameters to maintain connectivity between the Firewall Threat Defense device and the Firewall Management Center after deployment.

The following is a list of template validation checks done to ensure sanity of configurations that are required to manage the Firewall Threat Defense device using the data interface:

- You cannot apply a template in which manager access to the device is configured with the management interface to a device in which manager access to the device is configured with the data interface.
- You cannot apply a template in which manager access to the device is configured with the data interface to a device in which manager access to the device is configured with the management interface.
- You cannot apply a template in which manager access to the device is configured with the single WAN data interface to a device in which manager access to the device is configured with the dual WAN data interface.
- If any of the connectivity parameters do not match, you cannot apply a template in which manager access to the device is configured with the data interface to a device in which manager access to the device is configured with the data interface.

Perform the procedure given below to configure the template to manage Firewall Threat Defense devices using the data interface.

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the **Edit** (✎) icon of the template that you want to configure to manage Firewall Threat Defense devices using the data interface.

Step 3 Click the **Template Settings** tab.

Step 4 In the **General** tile, toggle the **Manage device by Data Interface** button.

Step 5 You will see a popup asking you to pick a data interface for manager access. Click **OK**.

Step 6 Click the **Interfaces** tab.

Step 7 Click the **Edit** icon of the data interface that you want to use for manager access. The first data interface – Ethernet1/1 (outside interface), is the data interface that is most commonly used for manager access.

Step 8 In the **Edit Physical Interface** window, click the **Manager Access** tab.

Step 9 Check the **Enable management access** checkbox.

Step 10 Click **OK**. You will see that the interface that you selected for manager access has been marked with **Manager Access**.

Step 11 Click the **DHCP** tab.

Step 12 Click the **DDNS Update Methods** tab.

Step 13 Click **+Add** to add a DDNS update method.

Step 14 In the **Add DDNS Update Method** window, enter a **Method Name** and choose **FMC only**.

Step 15 Set the **Update Interval** as per your requirement.

Step 16 Click **OK**. You will see the method that you created in the **DDNS Update Methods** table.

Step 17 Click the **DDNS Interface Settings** tab.

Step 18 Click **+Add** to add dynamic DNS configuration.

Step 19 In the **Add Dynamic DNS** configuration window, choose values for the following fields:

- **Interface** – Choose the interface enabled for manager access
- **Method Name** – Choose the method that you created.
- **Host Name** – Choose a variable for the hostname.

Do not edit the rest of the fields in this window.

Step 20 Click **OK**. The **DDNS Interface Settings** table is populated with the entry that you created.

Step 21 To configure the model mapping to ensure that the data interface set for manager access in the template matches the data interface selected for manager access on the device, click the **Template Settings** tab and click **Model Mapping**.

Step 22 Click **Add Model Mapping**.

Step 23 Choose the **Device Model** from the drop-down list.

Step 24 Map the date interface that is set for manager access in the template to the appropriate data interface on the device by choosing the interface from the **Model Interface** drop-down list.

Step 25 Click **Save**. The interface mappings are listed along with the device model and mapping status on the **Model Mapping** window. You can now apply the template on a device that is managed using the data interface.

Use Templates with Devices

When you register a device with the Firewall Management Center, you can choose a compatible template for the initial configuration. You can also apply a template to devices that are already registered with the Firewall Management Center.

Add a Device to the Management Center Using a Device Template

You can add a device to the Firewall Management Center using a device template with of the following options:

- Add a Device Using a Registration Key—Device Template, on page 36
- Add Devices Using Serial Numbers (Zero-Touch Provisioning)—Device Template, on page 47

Note Any change management tickets related to template configurations must be approved for the corresponding changes to be incorporated into the template application workflow. Only approved template configurations are used during template application.

Apply Templates to Existing Devices

You can apply or reapply a template to existing devices.

Apply a Template

You can apply a template to devices that are already registered with the Firewall Management Center. Application of a template on a device clears existing configurations and applies configurations from the template. However, the Firewall Threat Defense HA failover configurations are not cleared.

Application of a template changes device configurations only on the Firewall Management Center. You must explicitly deploy these device configuration changes to the Firewall Threat Defense device. You cannot roll back the applied configuration changes. However, you can apply another template with the required configurations.

Note Any change management tickets related to template configurations must be approved for the corresponding changes to be incorporated into the template application workflow. Only approved template configurations are used during template application.

Perform the procedure given below to apply the template on existing devices.

Procedure

Step 1 To apply the template from the **Template Management** window, choose **Devices** > **Template Management**.

a) Click the **More** (⋮) icon next to the template that you want to apply, and click **Apply**.
b) From the **Device** dropdown list, choose the **Device** on which you want to apply the template.
c) Click **Confirm** to initiate application of template on the device.

Step 2 (Optional) To apply the template from the **Associated Devices** window, choose **Devices > Template Management**.

a) Click the **Edit** (✎) icon of the template that you want to apply to a device.
b) Click **Associated Devices**.
c) In the **Associated Devices** window, click **Apply Template**.
d) From the **Device** drop-down list, choose the **Device** on which you want to apply the template.
e) Enter values for the **Variables** and **Network object overrides** fields.
f) Click **Apply** to initiate application of template on the device.

Reapply a Template

If you make any changes to the device or template that results in the configuration being out-of-sync, you can reapply the template to make the configuration in sync with the template.

Perform the procedure given below to reapply the template on a device.

Procedure

Step 1 Choose **Devices** > **Template Management**.

Step 2 Click the **Edit** (✎) icon of the template that you want to reapply to a device.

Step 3 Click **Associated Devices**.

Step 4 In the **Associated Devices** window, click **Reapply Template** for the device on which you want to reapply the template.

Note
If you want to reapply the template on all the associated devices in the template, click **Bulk Reapply** and click **Confirm**.

Step 5 On the **Reapply template** window, you can reuse the autopopulated **Variables** and **Network object overrides** values or enter new values.

Step 6 Click **Confirm** to initiate reapplication of the template on the device.

Monitoring Device Templates

You can monitor and verify the application of templates by viewing the devices listed in the **Associated Devices** window and by viewing the **Template Apply Report**.

View Associated Devices

The devices that are associated with the templates are listed in the **Associated Devices** window. Each device row displays the **Device Name, Sync Status, Template Application Status**, and **Applied Date**. You can also click **Reapply template** to reapply the template. Click the **Variable Summary** icon to display a summary of the variables in the template and the **Report** ( ) icon to download the Device Template Apply report. Click the **Delete** ( ) icon to remove the template from the device.

Click **Apply Template** if you want to apply a template on a device from the **Associated Devices** window. If you want to reapply the template on all the associated devices in the template, click **Bulk Reapply** and click **Confirm**.

The **Sync Status** can be either **Sync** or **Out-of-Sync**. If the status is displayed as **Sync**, it indicates that the template and device configurations are the same or in sync. If the status is displayed as **Out-of-Sync**, it indicates that there has been a change in configuration either on the device or in the template since the last time that the template was applied..

The association of the device with the template is not altered by the following conditions:

- Pending configuration changes on the device – The **Sync Status** does not change if there are pending configuration changes that have to applied on the device.
- Deployment of pending configuration changes on the device – The **Sync Status** does not change after deployment of pending configuration changes on the device.

The table below shows the **Sync** and **Out-of-Sync** scenarios that may occur.

Device Configurations Modified After Application of Template on Device	Template Configurations Modified After Application of Template on Device	Association Status
No	No	In Sync
Yes	No	Out of Sync
No	Yes	Out of Sync
Yes	Yes	Out of Sync

Generate a Template Apply Report

A **Template Apply Report** PDF is generated after the task to apply the template is completed. This report is generated on both successful and unsuccessful application of the template on the device. You will see a link to this report in the **Notifications > Tasks** window.

The **Template Apply Report** contains the following details:

- Template name
- Device model name
- Domain from which the template was applied
- Start and end time
- Status of the application of the template on the device

- Interface mapping information
- Variable values

There may be some configurations on the template that are not applied to the device due to incompatible device model or version. The report also contains details about such configurations. The report also contains any errors that are encountered when the application of the template fails. Application of a template on a device may fail due to any of the following reasons:

- Model mapping does not exist for the device model that is used.
- Values used for variables and network object overrides do not conform to routing policy or interface configuration rules. For example, the same IPv4 address has been used for two IPv4 address interface variables.
- Device or template is locked due to some other task that is being executed, such as application or modification of the template.

Audit Logs

Logs related to application of the device template, configuration updates, device template creation, and deletion, are logged under audit logs. The device template audit logs are added to the log both at the start and at the end of the task to apply the template on the device.

An audit diff file is also generated that enables you to view configuration changes that have been done during application of the template on the device. Perform the procedure given below to view the diff file.

Procedure

Step 1 Choose **System** (⚙) > **Monitoring** > **Audit**.

Step 2 The device template logs are logged under the subsystem **Devices** > **Template Management**. Click the **diff** icon to open a new window that displays the configuration changes that have been done during the application of the template on the device.

Troubleshooting Device Templates

Initial Troubleshooting

For initial troubleshooting, we recommend looking at the information in the Template Apply Report and notifications that come up on the Firewall Management Center UI when you run into an error. The Firewall Management Center log files also contain detailed debugging and troubleshooting info.

Follow the procedure given below for initial troubleshooting.

1. Check the errors mentioned in the **Template Apply Report**. For more information, see Template Apply Report.
2. Review variable values and check for overlaps and incompatibilities.

3. Check model mappings to ensure if the correct model mappings exist. Delete or add mappings accordingly.
4. See the Firewall Management Center audit logs to find any other issues and resolve them.

Consider the following error scenario. In a device template, the inside interface is configured with a static IPv4 variable - *$insideIPv4*.

The BGP IPv4 address is configured with an IPv4 BGP neighbor.

An overlapping IPv4 address is configured for the BGP neighbor and an interface.

Due to the issues mentioned above, the application of the device template fails and an error is displayed.

To troubleshoot this error, identify the error from the notification displayed on the UI.

```
IP Address 192.168.10.1 same as ip address of interface - 'inside'(Ethernet1/1)
```

Check the **Template Apply Report** for more information.

Enter correct values for the variables and apply the template again to ensure successful application of the template on the device.

Troubleshoot Device Registration

- Issue: Admin Password is incorrect or not provided during registration

  Scenario: If the admin password is not set on the device and if you have not provided the admin password during registration, the Firewall Threat Defense device provisioning will fail. In such a scenario, a *Provision Error* along with an **Enter Password** link is displayed.

  Workaround: Click **Enter Password** to enter a new password and click **Save**. Click **Confirm and Proceed** to trigger the onboarding again.

- If the admin password is already set on the device and you provide another admin password during registration, device provisioning will fail.

- Issue: Device registration in Firewall Management Center fails

  Workaround: Follow existing device registration troubleshooting steps. For more information, see Configure, Verify, and Troubleshoot Firepower Device Registration.

- Issue: Bulk Registration Request Fails in Firewall Management Center

  Scenario: The bulk registration request can fail due to a few scenarios:

  - You do not have the required permissions to perform template-related operations
  - Template is not visible from the request domain
  - Invalid CSV file provided

  Workaround: You can see logs for these errors in the VMS Shared and USM Shared log files. Fix the errors and initiate registration again.

- Issue: Device provisioning fails in Security Cloud Control due to some generic errors, such as communication with the device fails

  Workaround: Click Retry in the Provision Error to trigger the onboarding in Security Cloud Control again. You can also see the Security Cloud Control workflows for more information on the error and troubleshooting information.

Troubleshoot Cisco Security Cloud Integration

Issue: Cisco Security cloud integration not successful

Workaround: Follow Cisco Security Cloud integration troubleshooting steps. For more information, see Cisco Security Cloud Integration.

Troubleshoot Device Template Configuration Issues

Issue: Device template misconfigurations causing deployment failures after registration

Workaround: Follow the steps given below for initial troubleshooting.

1. Check the errors mentioned in the Template Apply Report.
2. Review variable values and check for overlaps and incompatibilities.
3. Check model mappings to ensure if the correct model mappings exist. Delete or add mappings accordingly.
4. See the Firewall Management Center audit logs to find any other issues and resolve them.

Troubleshoot Security Cloud Control Issues

- Issue: Device with serial number already claimed

  Workaround: Verify serial number and reinitiate onboarding.

- Issue: Security Cloud Control fails to claim devices

  Workaround: Select the device in the Security Cloud Control **Security Devices** window for more details on the error. You can see logs related to device claim issues in the VMS Shared and USM Shared log files. Click **Retry** to initiate registration again.

- Issue: Communication failures between Firewall Management Center and Security Cloud Control

  Scenario: Communication failures between Firewall Management Center and Security Cloud Control can cause failures during the Zero-Touch Provisioning (ZTP) device registration request.

  Workaround: Refresh the ZTP device status, retry ZTP registration, and delete the ZTP device. You can see logs regarding communication failure between the Firewall Management Center and Security Cloud Control in the Auth Daemon logs. For operational failures related to ZTP, you can see the logs in the VMS Shared and USM Shared log files.

History for Device Management using Device Templates

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Device management using device templates	7.6.0	7.4.1	Device templates enable deployment of multiple branch devices with pre-provisioned initial device configurations. You can use device templates to perform bulk zero-touch provisioning of multiple devices, apply day 2 configuration changes to multiple devices with different interface configurations, and clone configuration parameters from existing devices. New/modified screens: • **Devices** > **Device Management** > **Add** > **Device (Wizard)** • **Devices** > **Template Management** > **Add Device Template** • **Devices** > **Template Management** > **Add Model Mapping** • **Devices** > **Template Management** > **Edit a template** > **Template Settings** • **Integration** > **Cisco Security Cloud**

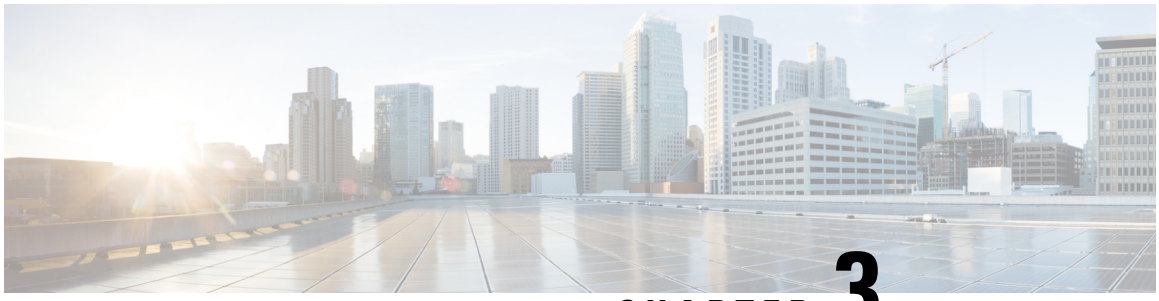

CHAPTER 3

Device Settings

After you add a device, you can edit device-related settings on the **Device** page.

1. Choose **Devices** > **Device Management**.
2. Next to the device you want to modify, click **Edit** ( ).
3. Click **Device**.

- Edit General Settings, on page 127
- Edit License Settings, on page 141
- View System Information, on page 141
- View the Inspection Engine, on page 142
- Edit Health Settings, on page 143
- Edit Management Settings, on page 153
- View Inventory Details, on page 195
- Edit Applied Policies, on page 196
- Edit Advanced Settings, on page 197
- Edit Deployment Settings, on page 201
- Edit Cluster Health Monitor Settings, on page 204
- History for Device Settings, on page 209

Edit General Settings

The **General** section of the **Device** page displays the settings described in the table below.

Figure 36: General

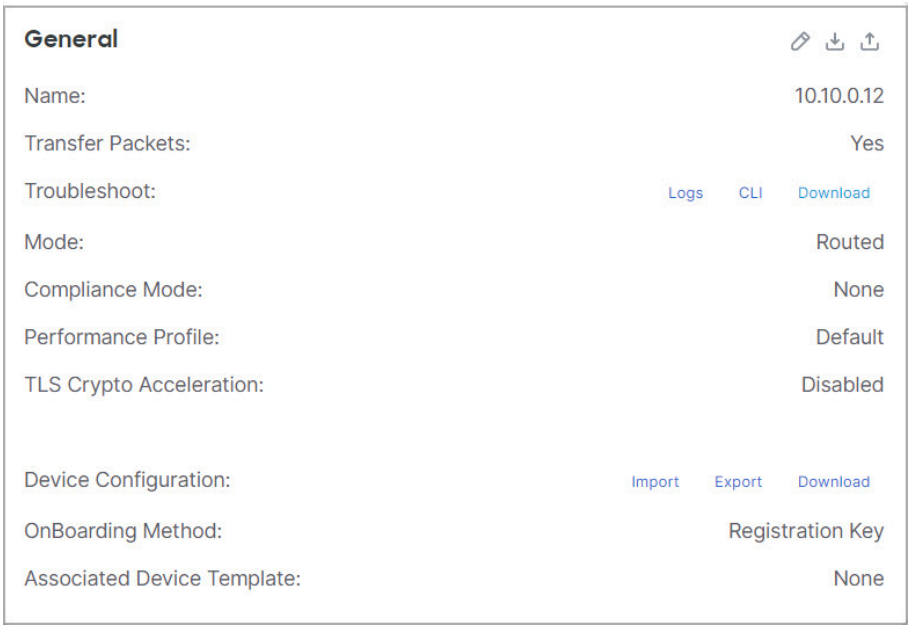

Table 4: General Section Table Fields

Field	Description
Name	The display name of the device on the Firewall Management Center.
Transfer Packets	This displays whether or not the managed device sends packet data with the events to the Firewall Management Center.
Troubleshoot	Lets you generate and download troubleshooting files and also see CLI command output. See Generate Troubleshooting Files, on page 129 and View CLI Output, on page 132.
Mode	The displays the mode of the management interface for the device: **routed** or **transparent**.
Compliance Mode	This displays the security certifications compliance for a device. Valid values are CC, UCAPL and None.
Performance Profile	This displays the core allocation performance profile for the device, as configured in the platform settings policy.
TLS Crypto Acceleration:	Shows whether TLS crypto acceleration is enabled or disabled.
Device Configuration	Lets you copy, export, or import a configuration. See Copy a Configuration to Another Device, on page 134 and Export and Import the Device Configuration, on page 136.
OnBoarding Method	Shows whether the device was registered using a registration key or using the serial number (zero-touch provisioning).

You can edit some of these settings from this section.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click **Edit** ( ).

Step 3 Click **Device**.

Step 4 In the **General** section, click **Edit** ( ).

a) Enter a **Name** for the managed device.

b) Check **Transfer Packets** to allow packet data to be stored with events on the Firewall Management Center.

c) Click **Force Deploy** to force deployment of current policies and device configuration to the device.

Note
Force-deploy consumes more time than the regular deployment since it involves the complete generation of the policy rules to be deployed on the Firewall Threat Defense.

Step 5 For **Troubleshoot** actions, see Generate Troubleshooting Files, on page 129 and View CLI Output, on page 132.

Step 6 For **Device Configuration** actions, see Copy a Configuration to Another Device, on page 134 and Export and Import the Device Configuration, on page 136.

Step 7 Click **Deploy**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Generate Troubleshooting Files

You can generate and download troubleshooting files for each device and also for all cluster nodes. For a cluster, you can download all files as a single compressed file. You can also include cluster logs for the cluster for cluster nodes.

You can alternatively trigger file generation from the **Devices** > **Device Management** > **More** ( ) > **Troubleshoot Files** menu.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device or cluster you want to view, click **Edit** ( ).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Device** or **Cluster**.

Step 4 Generate logs for the device or for all cluster nodes.

a) On the **General** > **Troubleshoot** section, click **Logs**.

Figure 37: Logs

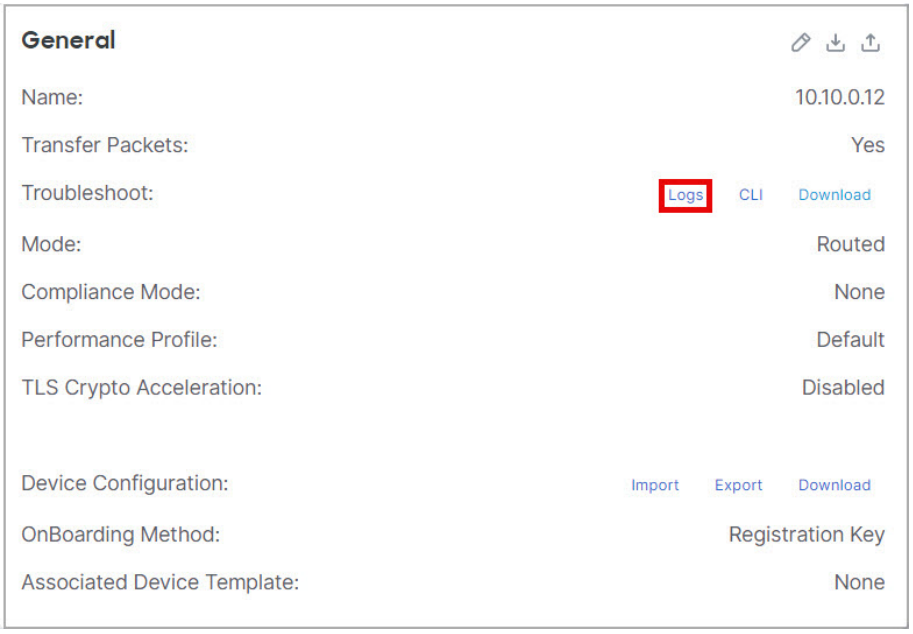

b) You are prompted to choose the logs you want to include. For a cluster, under **Device**, you can choose **All Devices** or an individual node. A cluster also has the **Cluster Logs** available.

Figure 38: Generate Troubleshoot Files

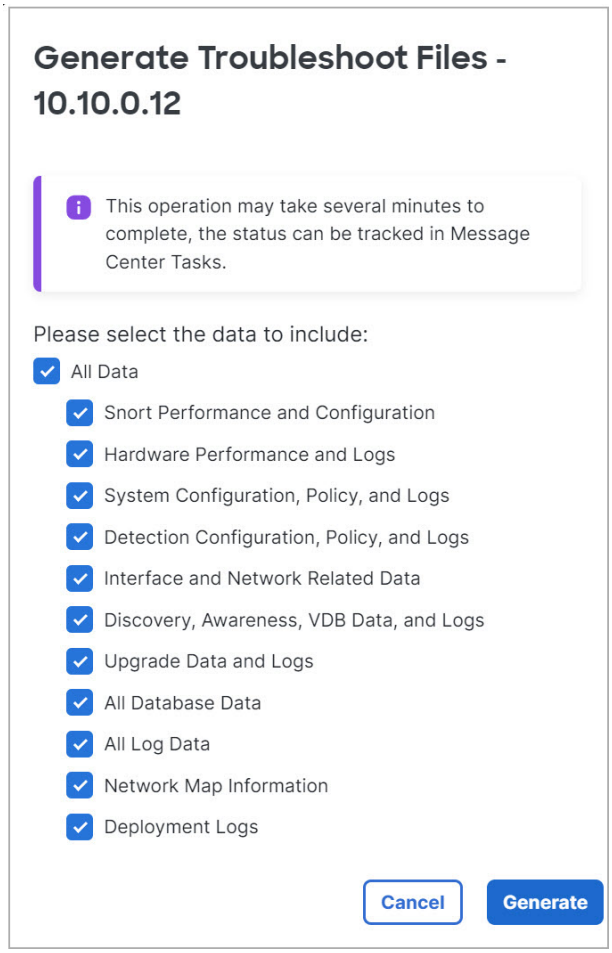

c) Click **Generate**.

Step 5 To download the generated logs, on the **General** > **Troubleshoot** section, click **Download**.

Figure 39: Download

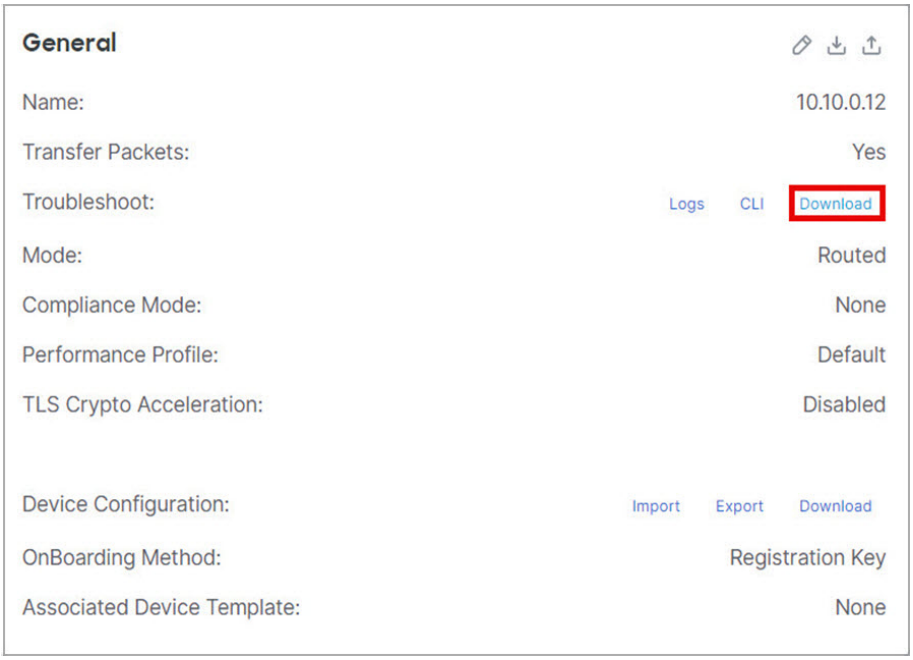

The logs are downloaded to your computer.

View CLI Output

You can view a set of pre-defined CLI outputs that can help you troubleshoot the device or cluster. You can also enter any **show** command and see the output.

For a device, the following commands are executed:

- **show version**
- **show asp drop**
- **show counters**
- **show int ip brief**
- **show blocks**
- **show cpu detailed**

For a cluster or cluster node:

- **show running-config cluster**
- **show cluster info**
- **show cluster info health**
- **show cluster info transport cp**

- **show version**
- **show asp drop**
- **show counters**
- **show arp**
- **show int ip brief**
- **show blocks**
- **show cpu detailed**
- **show interface** *ccl_interface*
- **ping** *ccl_ip* **size** *ccl_mtu* **repeat 2**

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device or cluster you want to view, click **Edit** (✎).

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Device** or **Cluster**.

Step 4 In the **General** > **Troubleshoot** section, click **CLI**.

Figure 40: CLI

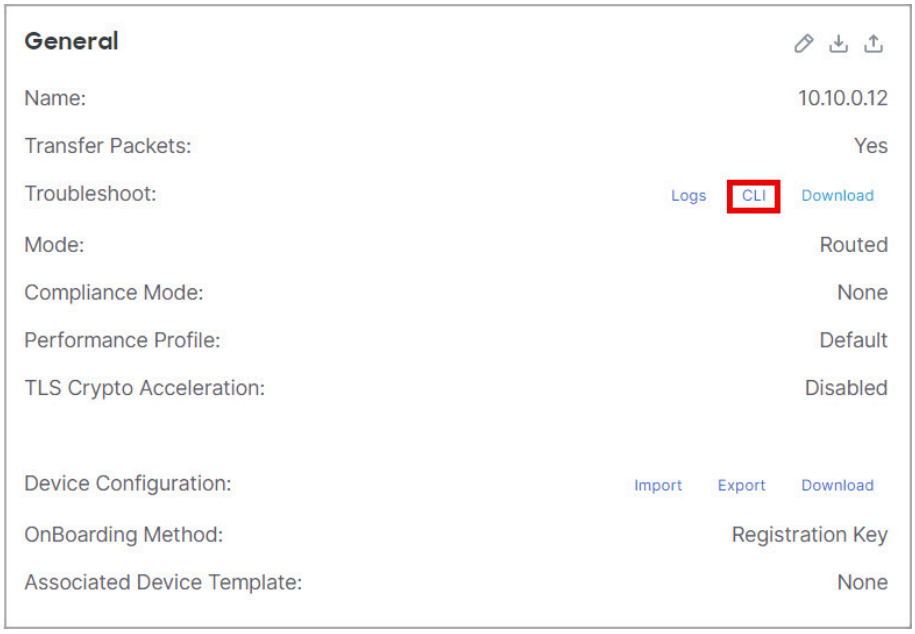

The **CLI Troubleshoot** dialog box appears with the pre-defined CLIs executed.

Figure 41: CLI Troubleshoot

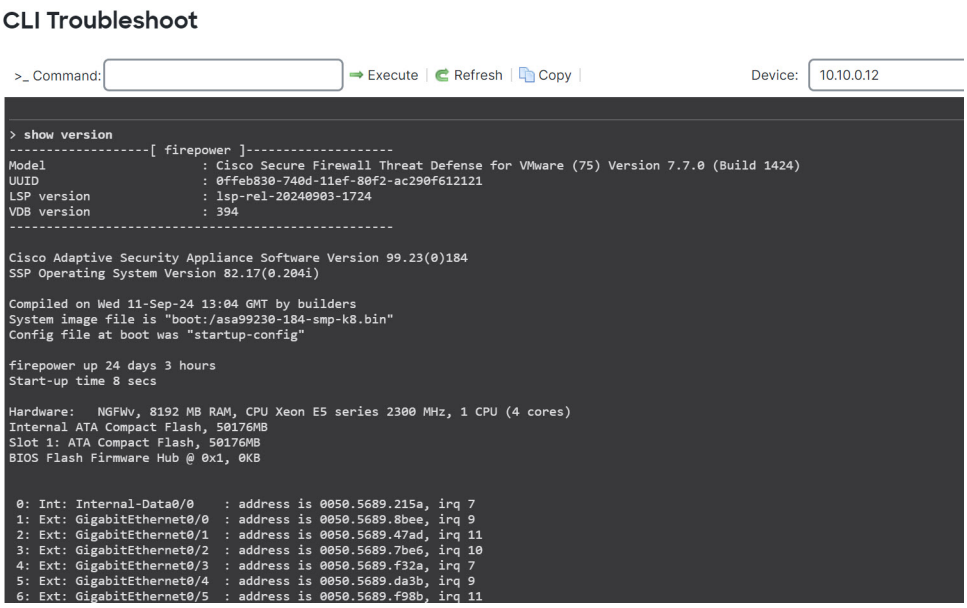

Step 5 On the **CLI Troubleshoot** dialog box, you can perform the following tasks.

- Enter a **show** command in the **Command** field, and click **Execute**. The new command output will be added to the window.
- Click **Refresh** to re-run the predefined CLIs.
- Click **Copy** to copy the output to your clipboard.
- For a cluster, choose a different node from the **Device** drop down list.

Step 6 Click **Close**.

Copy a Configuration to Another Device

When a new device is deployed in the network you can easily copy configurations and policies from a pre-configured device, instead of manually reconfiguring the new device.

Before you begin

Confirm that:

- The source and destination devices are the same model and are running the same version of the software.
- The source is either a standalone device or a high availability pair.
- The destination device is a standalone device.
- The source and destination devices have the same number of physical interfaces.
- The source and destination devices are in the same firewall mode: routed or transparent.

- The source and destination devices are in the same security-certifications-compliance mode.
- The source and destination devices are in the same domain.
- Configuration deployment is not in progress on either the source or the destination devices.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to modify, click **Edit** ( ).

Step 3 Click **Device**.

Step 4 In the **General** section, do one of the following:

Figure 42: Copy or Push Device Configuration

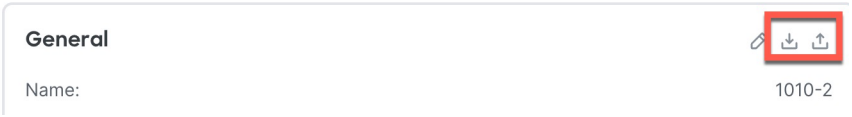

- Click **Get Device Configuration** ( ) to copy device configuration from another device to the new device. On the **Get Device Configuration** page, select the source device in the **Select Device** drop-down list.

Figure 43: Select Device

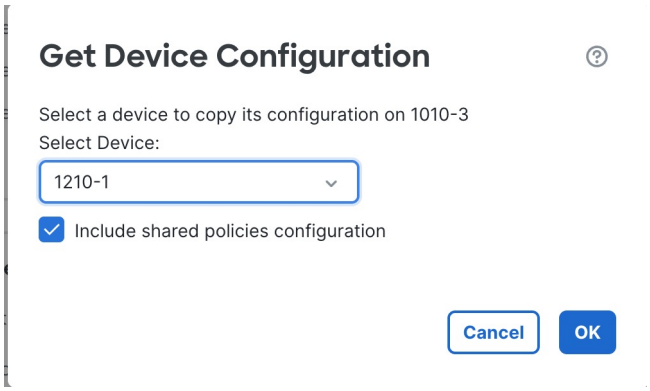

- Click **Push Device Configuration** ( ) to copy device configuration from the current device to the new device. On the **Push Device Configuration** page, select the destination to which configuration is to be copied in the **Target Device** drop-down list.

Step 5 (Optional) Check **Include shared policies configuration** check box to copy policies.

Shared policies like AC policy, NAT, Platform Settings and FlexConfig policies can be shared across multiple devices.

Step 6 Click **OK**.

You can monitor the status of the copy device configuration task on **Tasks** in the Message Center.

When the copy device configuration task is initiated, it erases the configuration on the target device and copies the configuration of the source device to the destination device.

Warning When you have completed the copy device configuration task, you cannot revert the target device to its original configuration.

Export and Import the Device Configuration

Note

- Export and import of device configuration between on-prem Firewall Management Center and Cloud-Delivered Firewall Management Center is not supported for shared policy and device policy.
- Export and import for the Cloud-Delivered Firewall Management Center is not supported for drop versions if underlying models are changed for policy in different drops.
- Export and import of device configuration is supported only if the device UUID, model and version are same.

You can export all of the the device-specific configuration configurable on the Device pages, including:

- Interfaces
- Inline Sets
- Routing
- DHCP
- VTEP
- Associated objects

You can then import the saved configuration for the same device in the following use cases:

- Moving the device to a different Firewall Management Center—First unregister the device from the original Firewall Management Center, then add the device to the new Firewall Management Center. Then you can import the saved configuration.
- Moving the device between domains—When you move a device between domains, some device-specific configuration is not retained because supporting objects (such as interface groups for security zones) do not exist in the new domain. By importing the configuration after the domain move, any necessary objects are created for that domain, and the device configuration is restored.
- Restore an old configuration—If you deployed changes that negatively impacted the operation of the device, you can import a backup copy of a known working configuration to restore a previous operational state.
- Reregistering a device—If you unregister a device from the Firewall Management Center, but then want to add it back, you can import the saved configuration.

See the following guidelines:

- You can only import the configuration to the same device (the UUID must match). You cannot import a configuration to a different device, even if it is the same model.
- Do not change the version running on the device between exporting and importing; the version must match.
- If you make inventory changes after your export (such as adding or deleting network modules or configuring or joining breakout ports), the device inventory will not match the Firewall Management Center. In this case, the device inventory will be maintained, and you will be prompted to sync the interfaces (see Sync Interface Changes with the Firewall Management Center, on page 789) and discard incompatible configuration in the Firewall Management Center when you try to deploy. You will have to repeat the inventory changes and related configuration in the Firewall Management Center.
- If you export a standalone configuration, you cannot import it to a high availability pair or vice versa.
- When moving the device to a different Firewall Management Center, the target Firewall Management Center version must be the same as the source version.
- If an object doesn't exist, it will be created. If an object exists, but the value is different, see below:

Table 5: Object Import Action

Scenario	Import Action
Object exists with the same name and value.	Reuse existing objects.
Object exists with the same name but different value.	Network and Port objects: Create object overrides for this device. See Object Overrides, on page 1361. Interface objects: Create new objects. For example, if both the type (security zone or interface group) and the interface type (routed or switched, for example) do not match, then a new object is created. All other objects: Reuse existing objects even though the values are different.
Object doesn't exist.	Create new objects.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device you want to edit, click **Edit** ( ).

Step 3 Click **Device**.

Step 4 Export the configuration.

a) In the **General** area, click **Export**.

Figure 44: Export Device Configuration

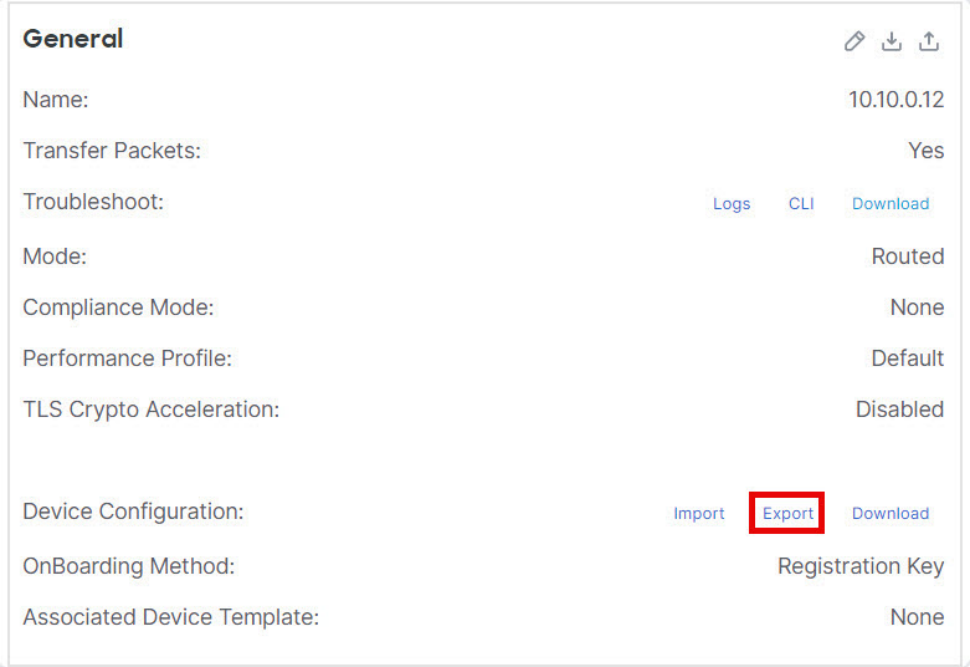

You are prompted to acknowledge the export; click **OK**.

Figure 45: Acknowledge Export

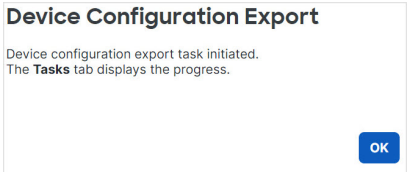

You can view the export progress in the **Tasks** page.

b) On the **Notifications** > **Tasks** page, ensure that the export has completed; click **Download Export Package**. Alternatively, you can click the **Download** button in the **General** area.

Figure 46: Export Task

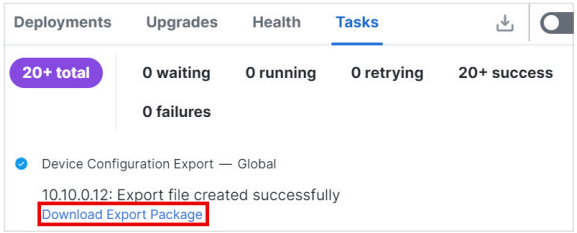

You are prompted to download the package; click **Click here to download the package** to save the file locally, and then click **OK** to exit the dialog box.

Figure 47: Download Package

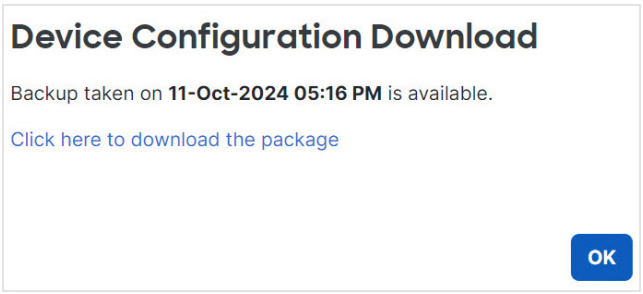

Step 5 Import the configuration.

a) In the **General** area, click **Import**.

Figure 48: Import Device Configuration

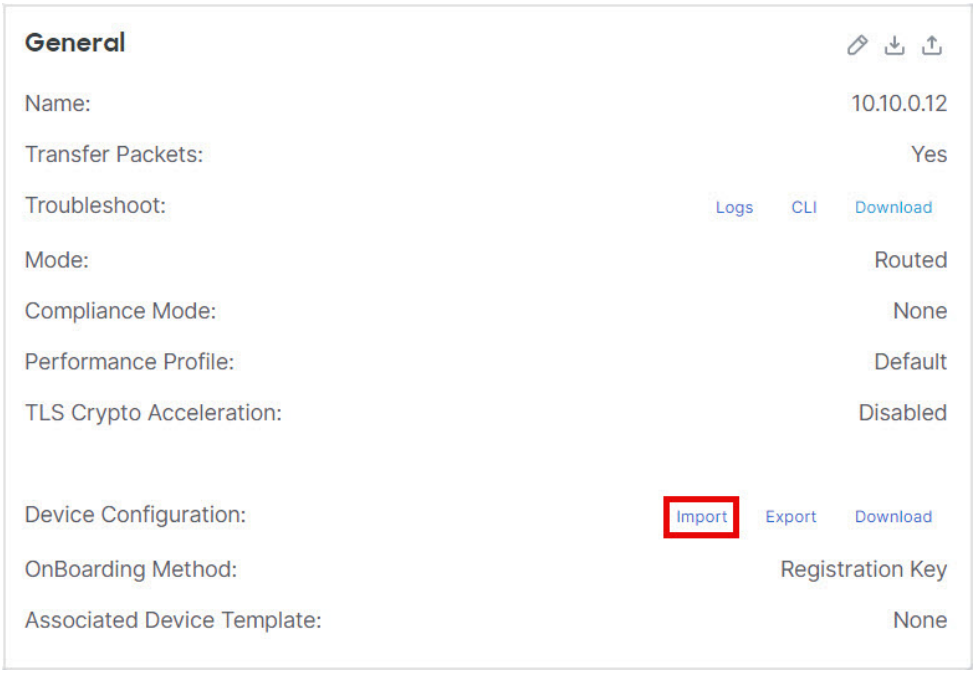

You are prompted to acknowledge that the current configuration will be replaced. Click **Yes** and then navigate to the configuration package (with the suffix .sfo; note that this file is different from the Backup/Restore files).

Figure 49: Import Package

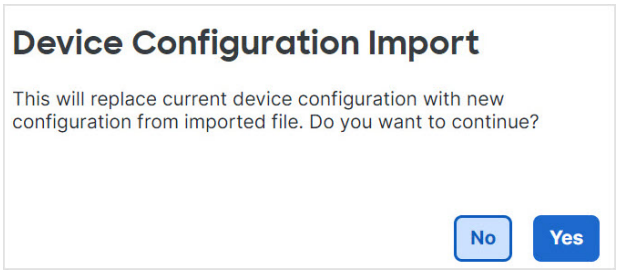

Figure 50: Navigate to Package

You are prompted to acknowledge the import; click **OK**.

Figure 51: Acknowledge Import

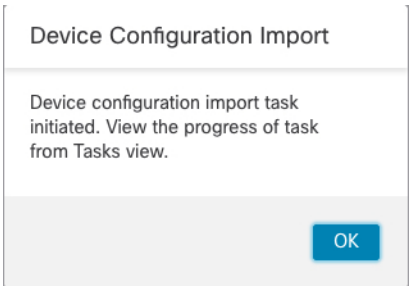

You can view the import progress in the **Tasks** page.

b) View the import reports so you can see what was imported. On the **Notifications** > **Tasks** page for the import task, click **View Import Report**.

Figure 52: View Import Report

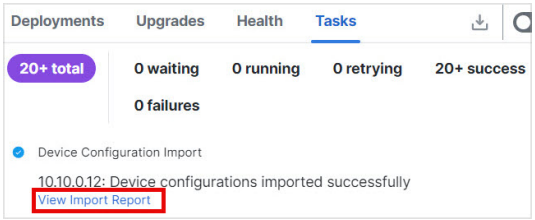

The **Device Configuration Import Reports** page provides links to available reports.

Device	Shared Policies	Device Configurations
0ffeb830-740d-11ef-80f2-ac290f612121	Report does not exist	Device configurations import report

c) Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Edit License Settings

The **License** section of the **Device** page displays the licenses enabled for the device.

You can enable licenses on your device if you have available licenses on your Firewall Management Center.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to enable or disable licenses, click **Edit** ( ).

Step 3 Click **Device**.

Step 4 In the **License** section, click **Edit** ( ).

Step 5 Check or clear the check box next to the license you want to enable or disable for the managed device.

Step 6 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

View System Information

The **System** section of the **Device** page displays a read-only table containing system information, as described in the following table.

You can also shut down or restart the device from this pane, using the icons at the top-right corner.

Figure 53: System

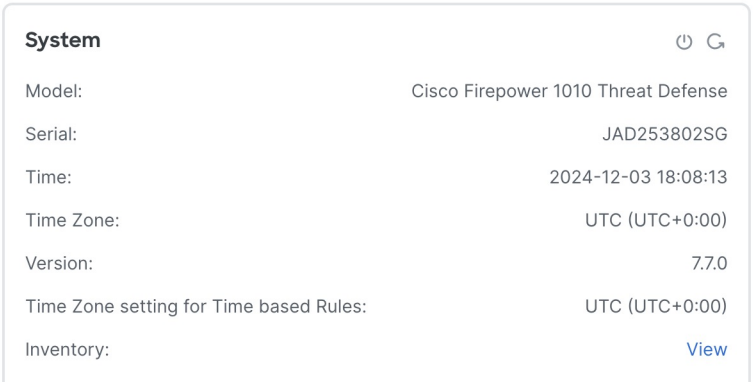

Table 6: System Section Table Fields

Field	Description
Shut Down Device ( )	Shuts down the device. See Shut Down or Restart the Device, on page 62.
Restart Device ( )	Restarts the device. See Shut Down or Restart the Device, on page 62.
Model	Model name and number of the managed device.
Serial	Serial number of the managed device's chassis.
Time	Current system time of the device.
Time Zone	Time zone.
Version	Version of the software currently installed on the managed device.
Time Zone setting for Time based rules	Current system time of the device in the time zone specified in device platform settings.
Inventory	Inventory details. See View Device Inventory.

View Device Inventory

Click **View** next to **Inventory** in the **System** section to view a table of device inventory information including Fans, Memory, CPU, Power Supply, Storage, and Network Modules.

The **Inventory Details** table displays information about all the Cisco products installed in the Firewall Threat Defense devices assigned with a product identifier (PID). The PID is the product name using which the product can be ordered.

Figure 54: Inventory Details

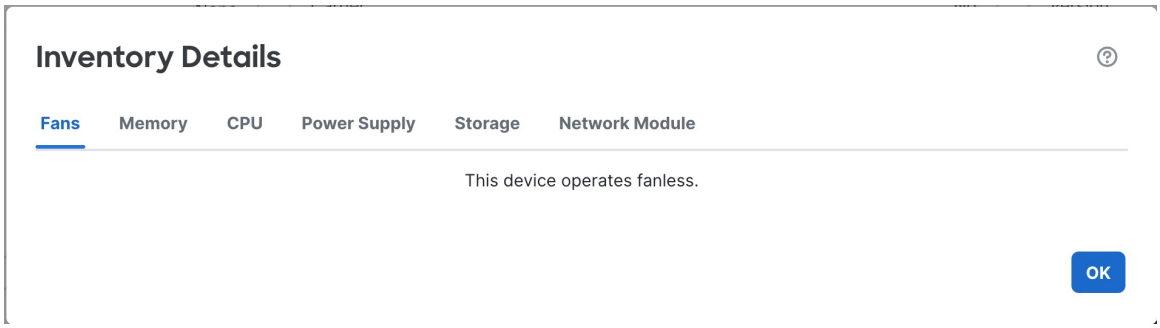

View the Inspection Engine

The Inspection Engine section of the **Device** page shows the inspection engine that is used on your device. Snort 3 is the only engine available for devices on version 7.7.

Edit Health Settings

The **Health** section of the **Device** page displays the information described in the table below.

Figure 55: Health

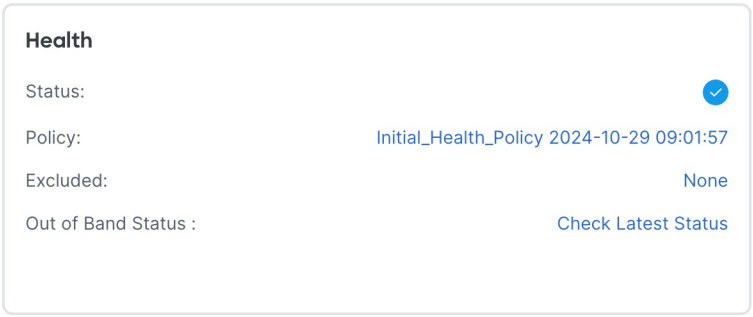

Table 7: Health Section Table Fields

Field	Description
Status	An icon that represents the current health status of the device. Clicking the icon displays the Health Monitor for the appliance.
Policy	A link to a read-only version of the health policy currently deployed at the device.
Excluded	A link to the **Health Exclude** page, where you can enable and disable health exclusion modules.
Out of Band Status	A link to the **Out-of-Band configuration details** dialog box where you can view out-of-band configuration changes made at the device CLI. You must acknowledge the configuration differences and manually match any changes you want to keep in the Firewall Management Center before the next deployment. See Out-of-Band Configuration Detection, on page 143.

Out-of-Band Configuration Detection

If you lose the management connection to your device, you can make select configuration changes directly at the device CLI to:

- Restore the management connection if you are using a data interface for manager access
- Make select configuration changes that can't wait until the connection is restored

Caution You are expected to know the commands that are required for recovery or emergency use. Do not use this feature to experiment with configuration changes. If you do not know which commands are required or are unsure about the effect of a command, we recommend that you contact Cisco TAC for guidance.

After the management connection is restored, the Firewall Management Center will detect the configuration changes on the device. It does not automatically update the device configuration in the Firewall Management Center; you must view the configuration differences, acknowledge that the device configuration is different, and then manually make the same changes in the Firewall Management Center before you deploy.

Caution When you deploy after acknowledgment, any configurations not present in the Firewall Management Center configuration will be overwritten on the device.

Guidelines for Out-of-Band Configuration

Supported Feature Areas in Recovery-Config Mode

You can configure the following feature areas at the diagnostic CLI in recovery-config mode:

- Interfaces
- Static Routes
- Dynamic Routing: BGP and OSPF
- Prefilters
- Site-to-site VPN

Like other diagnostic CLI commands, refer to the ASA command reference for more information about each command.

Unsupported Features

- Not supported in multi-instance mode.
- You cannot add or delete EtherChannels.
- Some platform-dependent interface commands such as speed, duplex, and shutdown are not supported.

High Availability and Clustering

- Recovery-config mode is only available on the active/control node.
- If a failover or cluster switchover occurs before you exit the recovery-config-mode session, the Firewall Management Center will not detect the change on the new active/control node. We recommend re-entering recovery-config mode on the new active/control node and making a small change to trigger discovery of all of your previous changes. Otherwise, if you do not manually match the changes in the Firewall Management Center, they will be overwritten at deployment without any notification.
- If you make out-of-band-configuration changes on the active/control node, but then, prior to a configuration sync, the high availability/cluster ends up in "split brain" mode (where multiple nodes become active/control because of a failover/cluster-control-link failure), then when the high availability/cluster returns to a healthy state, and a different node becomes active/control, then the configuration changes will be lost.
- If you have an active recovery-config-mode session, then new nodes cannot join or rejoin the high availability/cluster until the session is exited.

Additional Guidelines

- To modify an existing rule or route, you should delete the existing command using the **no** form of the command and then re-add the modified rule. This method avoids conflicts and errors. For example:

  Incorrect:

```
firepower# show running-config route
route outside 10.0.0.0 255.0.0.0 20.1.1.1 1
firepower# configure recovery-config

 CAUTION: The config CLI is for emergency use only. Use the config CLI if the management
 center is
unreachable, and use it only under exceptional circumstances, such as loss of connectivity
 or
to restore manager access. Do not change management center's auto-generated
configurations.

 After your management center is reachable, manually make the same configuration changes
 in the
management center. The management center cannot implement them automatically. When you
 deploy
from the management center, out-of-band configuration changes will be overwritten. Also,
 node join
will be blocked till config CLI session is active, so make sure to exit from the config
 CLI after
changes are made.

Would you like to proceed ? [Y]es/[N]o: y
firepower(recovery-config)# route outside 10.0.0.0 255.0.0.0 30.1.1.1
firepower(recovery-config)# exit
Unsaved changes are not kept if you reboot.Save changes to memory ? [Y]es/[N]o: y
Cryptochecksum: ccfc11a8 4e46d55e 0c99b5ae 3b18a8f1

3939 bytes copied in 0.70 secs
firepower# show running-config route
route outside 10.0.0.0 255.0.0.0 20.1.1.1 1
route outside 10.0.0.0 255.0.0.0 30.1.1.1 1
firepower#
```

  In this case, a second route is added instead of replacing the first route.

  Correct:

```
firepower# show running-config route
route outside 10.0.0.0 255.0.0.0 20.1.1.1 1
firepower# configure recovery-config

 CAUTION: The config CLI is for emergency use only. Use the config CLI if the management
 center is
unreachable, and use it only under exceptional circumstances, such as loss of connectivity
 or
to restore manager access. Do not change management center's auto-generated
configurations.

 After your management center is reachable, manually make the same configuration changes
 in the
management center. The management center cannot implement them automatically. When you
 deploy
from the management center, out-of-band configuration changes will be overwritten. Also,
 node join
will be blocked till config CLI session is active, so make sure to exit from the config
```

```
 CLI after
changes are made.

Would you like to proceed ? [Y]es/[N]o: y
firepower(recovery-config)# no route outside 10.0.0.0 255.0.0.0 20.1.1.1
firepower(recovery-config)# route outside 10.0.0.0 255.0.0.0 30.1.1.1
firepower(recovery-config)# exit
Unsaved changes are not kept if you reboot.Save changes to memory ? [Y]es/[N]o: y
Cryptochecksum: 81bcc51d 43771bbd 15b6dde6 afeb3442

3945 bytes copied in 0.70 secs
firepower# show running-config route
route outside 10.0.0.0 255.0.0.0 30.1.1.1 1
firepower#
```

- If you have auto rollback enabled (see Edit Deployment Settings, on page 201), and you lose management connectivity because of a deployment, you should not start an out-of-band configuration. Instead, either wait 20 minutes for auto rollback to the previous deployment to occur or manually roll back at the CLI using the **configure policy rollback** command (see Manually Roll Back the Configuration if the Firewall Management Center Loses Connectivity, on page 189). Auto rollback will overwrite out-of-band configuration changes if the management connection is still down.
- For prefilter rules, we don't recommend adding completely new rules (the **access-control advanced** command); integration of prefilter rules with the intrusion policy and logging requires the Firewall Management Center, which generates the rule ID and integrates it with other policies.
- All recovery-config-mode sessions will be logged in syslog with the username “enable_15”.

Access Recovery-Config Mode in the Diagnostic CLI

You can use the diagnostic CLI recovery-config mode to make out-of-band configuration changes when the management connection is down. Be sure to make the same changes in the Firewall Management Center; local changes will always be overwritten by the Firewall Management Center deployment.

For high availability and clustering, make your changes on the active/control node. This mode is not supported in multi-instance mode.

Procedure

Step 1 Connect to the device CLI using either the console port or SSH.

See Log Into the Command-Line Interface on the Device, on page 11.

Step 2 Access the diagnostic CLI.

system support diagnostic-cli

enable (Press enter without entering a password when prompted.)

Example:

```
> system support diagnostic-cli
firepower> enable
Password:
```

Step 3 Show the current running configuration for reference.

show runing-config

Note
You cannot enter **show** commands in recovery-config mode.

Step 4 Enter recovery-config mode.

configure recovery-config

Example:

```
firepower# configure recovery-config

 CAUTION: The config CLI is for emergency use only. Use the config CLI if the management
center is
unreachable, and use it only under exceptional circumstances, such as loss of connectivity
 or
to restore manager access. Do not change management center's auto-generated configurations.


 After your management center is reachable, manually make the same configuration changes
in the
management center. The management center cannot implement them automatically. When you
deploy
from the management center, out-of-band configuration changes will be overwritten. Also,
node join
will be blocked till config CLI session is active, so make sure to exit from the config CLI
 after
changes are made.

Would you like to proceed ? [Y]es/[N]o: y
firepower(recovery-config)#
```

Step 5 You can now enter select configuration commands.

Enter **?** to view available commands.

See Guidelines for Out-of-Band Configuration, on page 144 for supported feature areas.

See the ASA configuration guides or command reference for details about the commands.

Tip
Keep track of all of the commands you changed. Although the Firewall Management Center will show you the differential later, it's good practice to keep a record of your command changes in case you need to make iterative changes to restore the management connection.

Example:

```
firepower(recovery-config)# ?

  access-list           Configure an access control element
  as-path               BGP autonomous system path filter
  bfd                   BFD configuration commands
  bfd-template          BFD template configuration
  cluster               Cluster configuration
  community-list        Add a community list entry
  crypto                Configure IPSec, ISAKMP, Certification authority, key
  end                   Exit from configure mode
  exit                  Exit from config mode
  extcommunity-list     Add a extended community list entry
```

```
  group-policy          Configure or remove a group policy
  interface             Select an interface to configure
  ip                    Configure IP address pools
  ipsec                 Configure transform-set, IPSec SA lifetime and PMTU
                        Aging reset timer
  ipv6                  Configure IPv6 address pools
  ipv6                  Global IPv6 configuration commands
  isakmp                Configure ISAKMP options
  jumbo-frame           Configure jumbo-frame support
  management-interface  Management interface
  mtu                   Specify MTU(Maximum Transmission Unit) for an interface
  no                    Negate a command or set its defaults
  policy-list           Define IP Policy list
  prefix-list           Build a prefix list
  route                 Configure a static route for an interface
  route-map             Create route-map or enter route-map configuration mode
  router                Enable a routing process
  sla                   IP Service Level Agreement
  sysopt                Set system functional options
  tunnel-group          Create and manage the database of connection specific
                        records for IPSec connections
  vpdn                  Configure VPDN feature
  vrf                   Configure a VRF
  zone                  Create or show a Zone
firepower(recovery-config)#
```

Step 6 Exit recovery-config mode to be prompted to save your changes. Enter **exit** to exit each submode until you return to enable mode.

You can choose to save your changes to the startup configuration or keep changes only in the running configuration by not saving. Running configuration changes won't be retained after a reboot. If you make additional changes later and decide to save the configuration, all of your previous changes are also saved, since the entire running configuration is saved.

Deployment will be blocked while the recovery-config-mode session is open.

Example:

```
firepower(recovery-config)# interface Ethernet0/1
firepower(config-if)# ip address 10.0.0.2 255.0.0.0
firepower(config-if)# exit
firepower(recovery-config)# exit
Unsaved changes are not kept if you reboot. Save changes to memory ? [Y]es/[N]o: y

Cryptochecksum: 81a9073e f9535916 9c333d7e 9a3e5e76

3756 bytes copied in 0.70 secs
firepower#

Unsaved changes are not kept if you reboot. Save changes to memory ? [Y]es/[N]o:

Cryptochecksum: 81a9073e f9535916 9c333d7e 9a3e5e76

3756 bytes copied in 0.70 secs
firepower#
```

Step 7 Return to the Firewall Threat Defense CLI by typing Ctrl+a, then d, or you can enter **exit** to exit each mode.

Note
If you type Ctrl+a, then d to return to the Firewall Threat Defense CLI without first exiting recovery-config mode, the recovery-config-mode session will remain open, and deployment will be blocked.

Example:

```
firepower# exit

Logoff

User enable_1 logged in to firepower
Logins over the last 1 days: 4.  Last login: 20:42:51 UTC Dec 4 2024 from console
Failed logins since the last login: 0.
Type help or '?' for a list of available commands.
firepower> exit
Console connection detached.
>
```

Acknowledge the Out-of-Band Configuration

When the Firewall Management Center detects an out-of-band configuration change on a device, you must acknowledge the changes and match the configuration within the Firewall Management Center that you want to keep. Until you acknowledge the changes, deployment will be blocked.

Procedure

Step 1 Open the **Out-of-Band configuration details** dialog box.

Figure 56: Out-of-Band Configuration Details

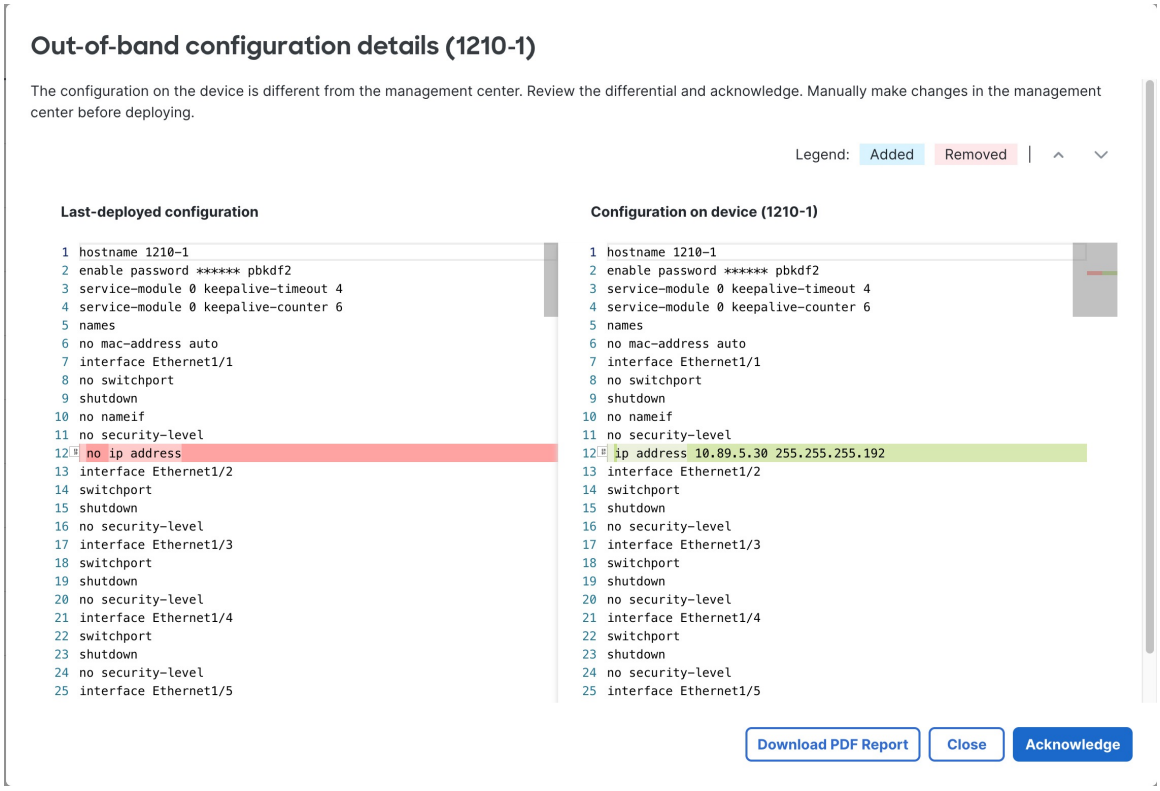

Note
Some commands, when set to a default setting, don't appear in the command output. However, the non-default command will show on either side as green (added) or red (removed). For example, if you add **no shutdown** to an interface in recovery-config mode, the **shutdown** command will show in red on the left **Last-deployed configuration** pane while **no shutdown** will *not* appear in the right **Configuration on device** pane. In this case, although the default setting for an interface is **shutdown**, the parser considers **no shutdown** to be the default and doesn't show it.

You can open the dialog box from multiple locations. For example, on the **Devices** > **Device Management** page, your device will have a warning. Click **View Details**.

Figure 57: Device Management Warning

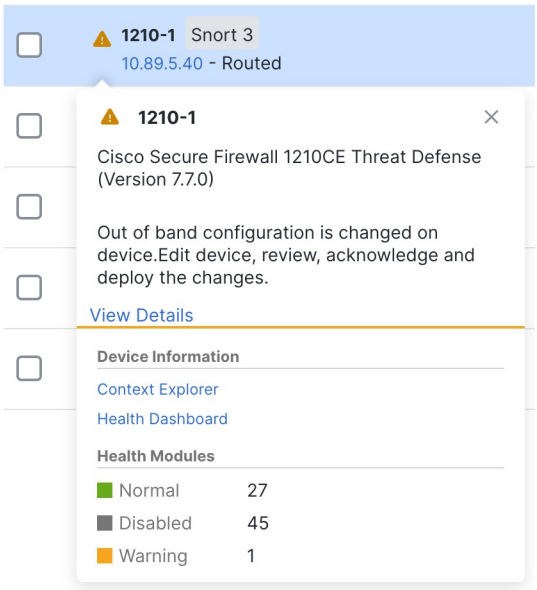

Or, from the **Devices** > **Device Management**, and then navigate to **Health** tile under **Device** tab, you can click **View Details**.

Figure 58: Health Out-of-Band Status

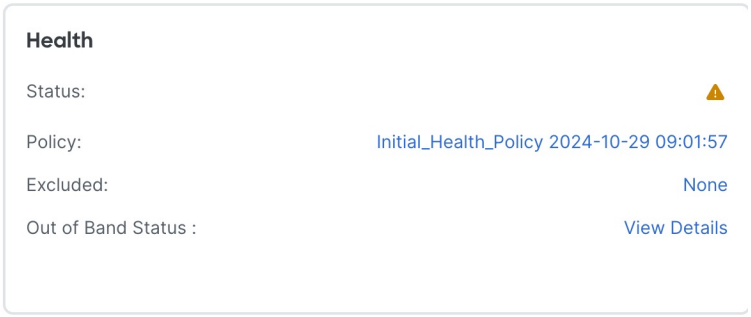

Note
If the out-of-band notification hasn't yet reached the Firewall Management Center, you can check for changes using the **Check Latest Status** link in the Health tile.

Step 2 Click **Download PDF Report** so you can refer to the configuration changes you need to make after you close the dialog box.

Or you can bring up the dialog box at any time to review the changes.

Step 3 Click **Acknowledge**, and then **Yes**.

Figure 59: Acknowledge

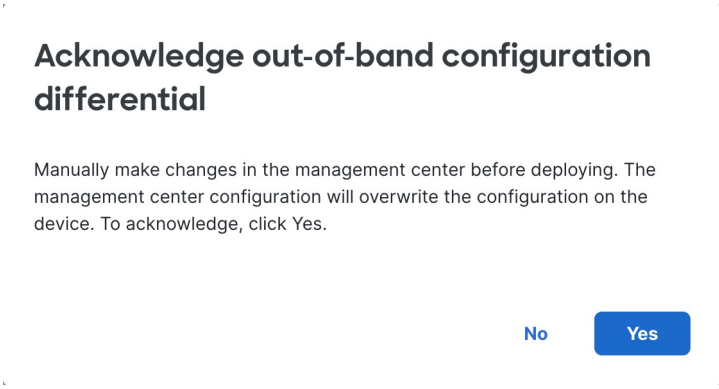

If you want to prevent an accidental deployment until after you've made your configuration changes, you can instead make the changes and then come back and click **Acknowledge**.

Step 4 Click **Close** on the **Out-of-Band configuration details** dialog box.

You can still revisit the dialog box to review the changes you need to make until you deploy. The status on the Device page changes to show you have acknowledged the out-of-band configuration:

Figure 60: Acknowledgement Status

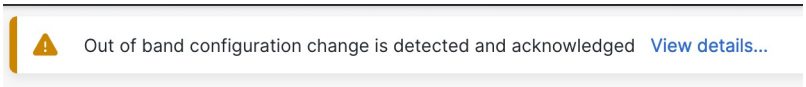

Step 5 Make the configuration changes that you made at the CLI.

You'll need to match the configuration CLI to Firewall Management Center screens; there aren't links from the CLI changes directly to screens.

If you don't want to keep your changes, you can simply deploy and overwrite the device configuration. You should make all necessary changes to maintain the management connection as well as any other changes you want to keep. For example, if you changed the IP address at the CLI, you need to go to the **Interfaces** page, edit the interface, and set that IP address to match:

Figure 61: Match the IP Address Change

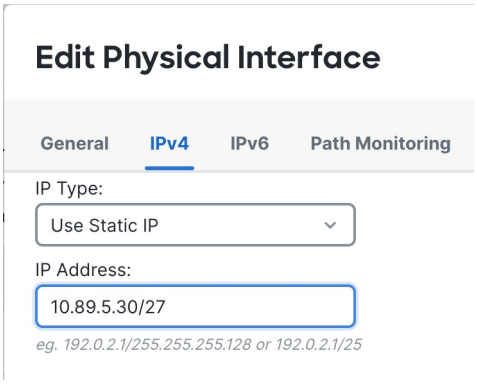

There is no checking mechanism that you made the same change; you could set the IP address differently if you want.

Step 6 Deploy configuration changes; see Deploy Configuration Changes, on page 259.

After you deploy, you can view the configuration differential—whether you made the changes or not—on the **System** (⚙) > **Monitoring** > **Audit** page. Check for the subsystem called **Device > Device Management > Out of band changes**.

Edit Management Settings

These settings control how the Firewall Management Center establishes the management connection with the device.

Configure a Redundant Manager Access Data Interface

When you use a data interface for manager access, you can configure a secondary data interface to take over management functions if the primary interface goes down. You can configure only one secondary interface. The device uses SLA monitoring to track the viability of the static routes and an ECMP zone that contains both interfaces so management traffic can use both interfaces.

Before you begin

- The secondary interface needs to be in a separate security zone from the primary interface.
- All of the same requirements apply to the secondary interface as apply to the primary interface. See Using the Firewall Threat Defense Data Interface for Management, on page 4.

Procedure

Step 1 On the **Devices** > **Device Management** page, click **Edit** (✎) for the device.

Step 2 Enable manager access for the secondary interface.

This setting is in addition to standard interface settings such as enabling the interface, setting the name, setting the security zone, and setting a static IPv4 address.

a) Choose **Interfaces** > **Edit Physical Interface** > **Manager Access**.
b) Check **Enable management on this interface for the Manager**.
c) Click **OK**.

Both interfaces show **(Manager Access)** in the interface listing.

Figure 62: Interface Listing

Interface	Logical Name	Type	Security Zones
Diagnostic1/1	diagnostic	Physical	
Ethernet1/1 (Manager Access)	outside	Physical	outside
Ethernet1/2		Physical	
Ethernet1/3		Physical	
Ethernet1/4		Physical	
Ethernet1/5		Physical	
Ethernet1/6		Physical	
Ethernet1/7		Physical	
Ethernet1/8 (Manager Access)	redundant	Physical	mgmt

Step 3 Add the secondary address to the **Management** settings.

a) Click **Device**, and view the **Management** area.

b) Click **Edit** (✎).

Figure 63: Edit Management Address

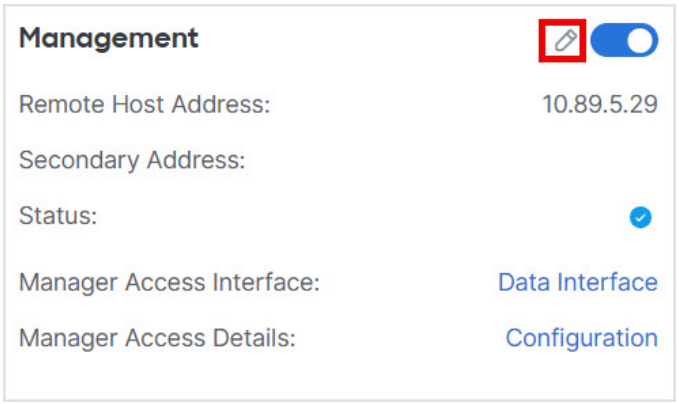

c) In the **Management** dialog box, modify the name or IP address in the **Secondary Address** field

Figure 64: Management IP Address

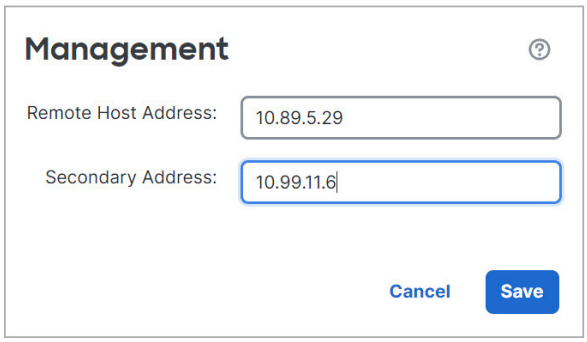

d) Click **Save**.

Step 4 Create an ECMP zone with both interfaces.

a) Click **Routing**.
b) From the virtual router drop-down, choose the virtual router in which the primary and secondary interfaces reside.
c) Click **ECMP**, and then click **Add**.
d) Enter a **Name** for the ECMP zone.
e) Select the primary and secondary interfaces under the **Available Interfaces** box, and then click **Add**.

Figure 65: Add an ECMP Zone

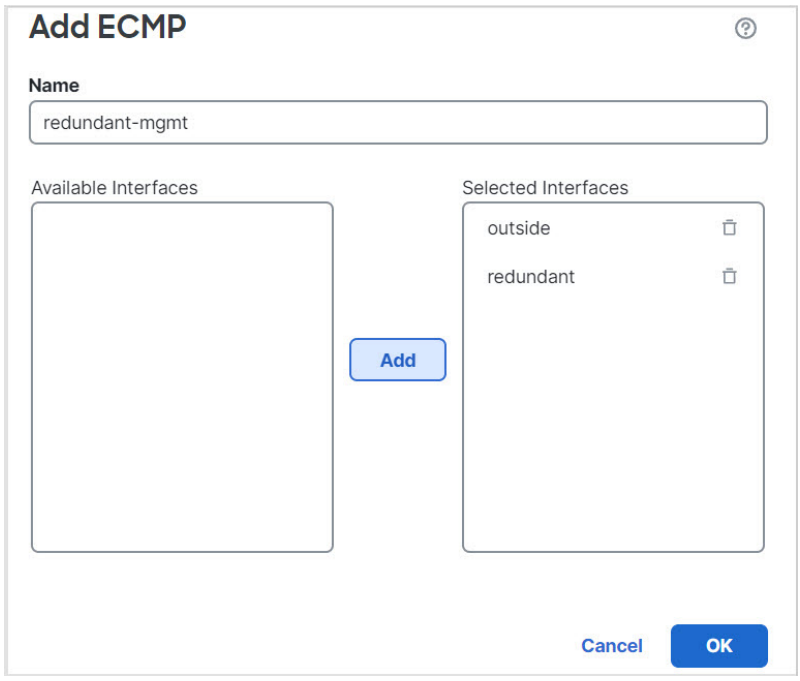

f) Click **OK**, and then **Save**.

Step 5 Add equal-cost default static routes for both interfaces and enable SLA tracking on both.

The routes should be identical except for the gateway and should both have metric 1. The primary interface should already have a default route that you can edit.

Figure 66: Add/Edit Static Route

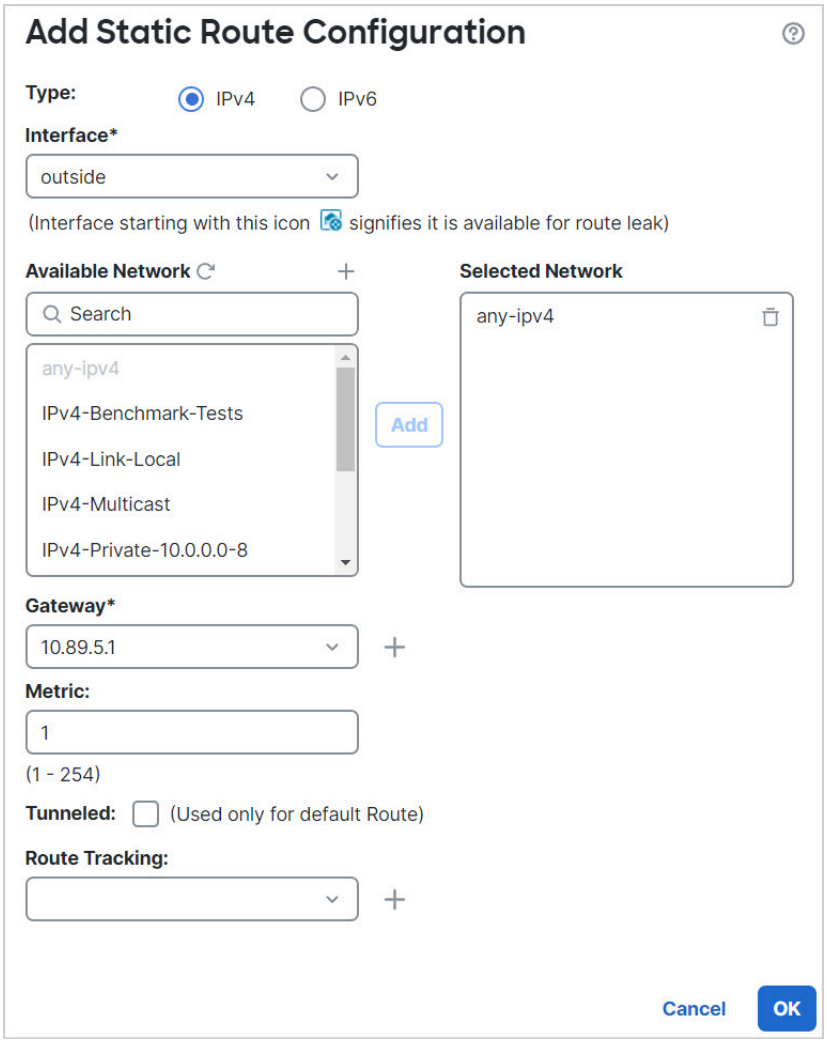

a) Click **Static Route**.
b) Either click **Add Route** to add a new route, or click **Edit** (✎) for an existing route.
c) From the **Interface** drop-down, choose the interface.
d) For the destination network, select **any-ipv4** from the **Available Networks** box and click **Add**.
e) Enter the default **Gateway**.
f) For **Route Tracking**, click **Add** (+) to add a new SLA monitor object.
g) Enter the required parameters including the following:

- The **Monitor Address** as the Firewall Management Center IP address.
- The zone for the primary or secondary management interface in **Available Zones**; for example, choose the outside zone for the primary interface object, and the mgmt zone for the secondary interface object.

See SLA Monitor, on page 1441 for more information.

Figure 67: Add SLA Monitor

h) Click **Save**, then choose the SLA object you just created in the **Route Tracking** drop-down list.
i) Click **OK**, and then **Save**.
j) Repeat for the default route for the other management interface.

Step 6 Deploy configuration changes; see Deploy Configuration Changes, on page 259.

As part of the deployment for this feature, the Firewall Management Center enables the secondary interface for management traffic, including auto-generated policy-based routing configuration for management traffic to get to the right data interface. The Firewall Management Center also deploys a second instance of the **configure network management-data-interface** command. Note that if you edit the secondary interface at the CLI, you cannot configure the gateway or otherwise alter the default route, because the static route for this interface can only be edited in the Firewall Management Center.

Change Manager Access Interface Settings

Changing any manager interface settings on the device or on the Firewall Management Center can disrupt the management connection. See the following scenarios to change interface settings and reestablish the management connection.

Change the Device IP Address

Change the device IP address, and then update the address in the Firewall Management Center.

Set the Device IP Address

Use one of the following methods to set the manager access interface IP address.

Modify Firewall Threat Defense Management Interfaces at the CLI

Modify the management interface settings on the managed device using the CLI. Many of these settings are ones that you set when you performed the initial setup; this procedure lets you change those settings, and set additional settings such as enabling an event interface if your model supports it, or adding static routes.

Note This topic applies to the dedicated Management interface. You can alternatively configure a data interface for management. If you want to change network settings for that interface, you should do so within Firewall Management Center and not at the CLI. If you need to troubleshoot a disrupted management connection, and need to make changes directly on the Firewall Threat Defense, see Modify the Firewall Threat Defense Data Interface Used for Management at the CLI, on page 164.

For information about the Firewall Threat Defense CLI, see the Cisco Secure Firewall Threat Defense Command Reference.

Note When using SSH, be careful when making changes to the management interface; if you cannot re-connect because of a configuration error, you will need to access the device console port.

Note If you change the device management IP address, then see the following tasks for Firewall Management Center connectivity depending on how you identified the Firewall Management Center during initial device setup using the **configure manager add** command (see Register With a New Management Center, on page 56):

- **IP address—No action.** If you identified the Firewall Management Center using a reachable IP address, then the management connection will be reestablished automatically after several minutes. We recommend that you also change the device IP address shown in Firewall Management Center to keep the information in sync; see Update the Hostname or IP Address in the Firewall Management Center, on page 170. This action can help the connection reestablish faster. **Note**: If you specified an unreachable Firewall Management Center IP address, then see the procedure for NAT ID below.
- **NAT ID only—Manually reestablish the connection.** If you identified the Firewall Management Center using only the NAT ID, then the connection cannot be automatically reestablished. In this case, change the device management IP address in Firewall Management Center according to Update the Hostname or IP Address in the Firewall Management Center, on page 170.

Note In a High Availability Firewall Management Center configuration, when you modify the management IP address from the device CLI or from the Firewall Management Center, the secondary Firewall Management Center does not reflect the changes even after an HA synchronization. To ensure that the secondary Firewall Management Center is also updated, switch roles between the two Firewall Management Centers, making the secondary Firewall Management Center the active unit. Modify the management IP address of the registered device on the device management page of the now active Firewall Management Center.

Before you begin

- You can create user accounts that can log into the CLI using the **configure user add** command; see Add an Internal User at the CLI, on page 215. You can also configure AAA users according to External Authentication, on page 952.

Procedure

Step 1 Connect to the device CLI, either from the console port or using SSH.

See Log Into the Command-Line Interface on the Device, on page 11.

Step 2 Log in with the Admin username and password.

Step 3 (Firepower 4100/9300/Secure Firewall 4200 only) Enable the second management interface as an event-only interface.

configure network management-interface enable management1

configure network management-interface disable-management-channel management1

You always need a management interface for management traffic. If your device has a second management interface, you can enable it for event-only traffic.

You can optionally disable events for the main management interface using the **configure network management-interface disable-events-channel** command. In either case, the device will try to send events on the event-only interface, and if that interface is down, it will send events on the management interface even if you disable the event channel.

You cannot disable both event and management channels on an interface.

To use a separate event interface, you also need to enable an event interface on the Firewall Management Center. See the Cisco Secure Firewall Management Center Administration Guide.

Example:

```
> configure network management-interface enable management1
Configuration updated successfully

> configure network management-interface disable-management-channel management1
Configuration updated successfully

>
```

Step 4 Configure the IP address of the management interface and/or event interface:

If you do not specify the *management_interface* argument, then you change the network settings for the default management interface. When configuring an event interface, be sure to specify the *management_interface* argument. The event interface can be on a separate network from the management interface, or on the same network. If you are connected to the interface you are configuring, you will be disconnected. You can re-connect to the new IP address.

a) Configure the IPv4 address:

- Manual configuration:

  configure network ipv4 manual *ip_address netmask gateway_ip* [*management_interface*]

  Note that the *gateway_ip* in this command is used to create the default route for the device. If you configure an event-only interface, then you must enter the *gateway_ip* as part of the command; however, this entry just configures the default route to the value you specify and does not create a separate static route for the eventing interface. If you are using an event-only interface on a different network from the management interface, we recommend that you set the *gateway_ip* for use with the management interface, and then create a static route separately for the event-only interface using the **configure network static-routes** command.

  Example:

  ```
  > configure network ipv4 manual 10.10.10.45 255.255.255.0 10.10.10.1 management1
  Setting IPv4 network configuration.
  Network settings changed.

  >
  ```

- DHCP (supported on the default management interface only):

  configure network ipv4 dhcp

b) Configure the IPv6 address:

- Stateless autoconfiguration:

configure network ipv6 router [*management_interface*]

Example:

```
> configure network ipv6 router management0
Setting IPv6 network configuration.
Network settings changed.

>
```

- Manual configuration:

  configure network ipv6 manual *ip6_address ip6_prefix_length* [*ip6_gateway_ip*] [*management_interface*]

  Note that the *ipv6_gateway_ip* in this command is used to create the default route for the device. If you configure an event-only interface, then you must enter the *ipv6_gateway_ip* as part of the command; however, this entry just configures the default route to the value you specify and does not create a separate static route for the eventing interface. If you are using an event-only interface on a different network from the management interface, we recommend that you set the *ipv6_gateway_ip* for use with the management interface, and then create a static route separately for the event-only interface using the **configure network static-routes** command.

  Example:

```
> configure network ipv6 manual 2001:0DB8:BA98::3210 64 management1
Setting IPv6 network configuration.
Network settings changed.

>
```

- DHCPv6 (supported on the default management interface only):

  configure network ipv6 dhcp

Step 5 For IPv6, enable or disable ICMPv6 Echo Replies and Destination Unreachable messages. These messages are enabled by default.

configure network ipv6 destination-unreachable {**enable** | **disable**}

configure network ipv6 echo-reply {**enable** | **disable**}

You might want to disable these packets to guard against potential denial of service attacks. Disabling Echo Reply packets means you cannot use IPv6 ping to the device management interfaces for testing purposes.

Example:

```
> configure network ipv6 destination-unreachable disable
> configure network ipv6 echo-reply disable
```

Step 6 Enable a DHCP server on the default management interface to provide IP addresses to connected hosts:

configure network ipv4 dhcp-server-enable *start_ip_address end_ip_address*

Example:

```
> configure network ipv4 dhcp-server-enable 10.10.10.200 10.10.10.254
```

```
DHCP Server Enabled

>
```

You can only configure a DHCP server when you set the management interface IP address manually. This command is not supported on the Firewall Management Center Virtual. To display the status of the DHCP server, enter **show network-dhcp-server**:

```
> show network-dhcp-server
DHCP Server Enabled
10.10.10.200-10.10.10.254
```

Step 7 Add a static route for the event-only interface if the Firewall Management Center is on a remote network; otherwise, all traffic will match the default route through the management interface.

configure network static-routes {**ipv4** | **ipv6**} **add** *management_interface destination_ip netmask_or_prefix gateway_ip*

For the *default* route, do not use this command; you can only change the default route gateway IP address when you use the **configure network ipv4** or **ipv6** commands (see Step 4, on page 160).

Example:

```
> configure network static-routes ipv4 add management1 192.168.6.0 255.255.255.0 10.10.10.1
Configuration updated successfully

> configure network static-routes ipv6 add management1 2001:0DB8:AA89::5110 64
2001:0DB8:BA98::3211
Configuration updated successfully

>
```

To display static routes, enter **show network-static-routes** (the default route is not shown):

```
> show network-static-routes
---------------[ IPv4 Static Routes ]---------------
Interface                 : management1
Destination               : 192.168.6.0
Gateway                   : 10.10.10.1
Netmask                   : 255.255.255.0
[…]
```

Step 8 Set the hostname:

configure network hostname *name*

Example:

```
> configure network hostname farscape1.cisco.com
```

Syslog messages do not reflect a new hostname until after a reboot.

Step 9 Set the search domains:

configure network dns searchdomains *domain_list*

Example:

```
> configure network dns searchdomains example.com,cisco.com
```

Set the search domain(s) for the device, separated by commas. These domains are added to hostnames when you do not specify a fully-qualified domain name in a command, for example, **ping system**. The domains are used only on the management interface, or for commands that go through the management interface.

Step 10 Set up to 3 DNS servers, separated by commas:

configure network dns servers *dns_ip_list*

Example:

```
> configure network dns servers 10.10.6.5,10.20.89.2,10.80.54.3
```

Step 11 Set the remote management port for communication with the Firewall Management Center:

configure network management-interface tcpport *number*

Example:

```
> configure network management-interface tcpport 8555
```

The Firewall Management Center and managed devices communicate using a two-way, TLS-1.3-encrypted communication channel, which by default is on port 8305.

Note
Cisco **strongly** recommends that you keep the default settings for the remote management port, but if the management port conflicts with other communications on your network, you can choose a different port. If you change the management port, you must change it for **all** devices in your deployment that need to communicate with each other.

Step 12 (Firewall Threat Defense only) Set the management or eventing interface MTU. The MTU is 1500 bytes by default.

configure network mtu [*bytes*] [*interface_id*]

- *bytes*—Sets the MTU in bytes. For the management interface, the value can be between 64 and 1500 if you enable IPv4, and 1280 to 1500 if you enable IPv6. For the eventing interface, the value can be between 64 and 9000 if you enable IPv4, and 1280 to 9000 if you enable IPv6. If you enable both IPv4 and IPv6, then the minimum is 1280. If you do not enter the *bytes*, you are prompted for a value.
- *interface_id*—Specifies the interface ID on which to set the MTU. Use the **show network** command to see available interface IDs, for example management0, management1, br1, and eth0, depending on the platform. If you do not specify an interface, then the management interface is used.

Example:

```
> configure network mtu 8192 management1
MTU set successfully to 1500 from 8192 for management1
Refreshing Network Config...
NetworkSettings::refreshNetworkConfig MTU value at start 8192

Interface management1 speed is set to '10000baseT/Full'
NetworkSettings::refreshNetworkConfig MTU value at end 8192
```

>

Step 13 Configure an HTTP proxy. The device is configured to directly-connect to the internet on ports TCP/443 (HTTPS) and TCP/80 (HTTP). You can use a proxy server, to which you can authenticate via HTTP Digest. After issuing the command, you are prompted for the HTTP proxy address and port, whether proxy authentication is required, and if it is required, the proxy username, proxy password, and confirmation of the proxy password.

Note
For proxy password on Firewall Threat Defense, you can use A-Z, a-z, and 0-9 characters only.

configure network http-proxy

Example:

```
> configure network http-proxy
Manual proxy configuration
Enter HTTP Proxy address: 10.100.10.10
Enter HTTP Proxy Port: 80
Use Proxy Authentication? (y/n) [n]: Y
Enter Proxy Username: proxyuser
Enter Proxy Password: proxypassword
Confirm Proxy Password: proxypassword
```

Step 14 If you change the device management IP address, then see the following tasks for Firewall Management Center connectivity depending on how you identified the Firewall Management Center during initial device setup using the **configure manager add** command (see Register With a New Management Center, on page 56):

- **IP address—No action.** If you identified the Firewall Management Center using a reachable IP address, then the management connection will be reestablished automatically after several minutes. We recommend that you also change the device IP address shown in Firewall Management Center to keep the information in sync; see Update the Hostname or IP Address in the Firewall Management Center, on page 170. This action can help the connection reestablish faster. **Note**: If you specified an unreachable Firewall Management Center IP address, then you must manually reestablish the connection using Update the Hostname or IP Address in the Firewall Management Center, on page 170.
- **NAT ID only—Manually reestablish the connection.** If you identified the Firewall Management Center using only the NAT ID, then the connection cannot be automatically reestablished. In this case, change the device management IP address in Firewall Management Center according to Update the Hostname or IP Address in the Firewall Management Center, on page 170.

Modify the Firewall Threat Defense Data Interface Used for Management at the CLI

If the management connection between the Firewall Threat Defense and the Firewall Management Center was disrupted, and you want to specify a new data interface to replace the old interface, use the Firewall Threat Defense CLI to configure the new interface.

If the management connection is active, then you should make any changes to an existing data interface using the Firewall Management Center (see Modify the Firewall Threat Defense Data Interface Used for Management in the GUI, on page 167). For initial setup of the data management interface, see the **configure network management-data-interface** command in Complete the Firewall Threat Defense Initial Configuration Using the CLI, on page 19.

For high-availability pairs, perform all CLI steps on both units. Within the Firewall Management Center, perform steps only on the active unit. Once the configuration changes are deployed, the standby unit synchronizes configuration and other state information from the active unit.

Note This topic applies to the data interface that you configured for Management, not the dedicated Management interface. If you want to change network settings for the Management interface, see Modify Firewall Threat Defense Management Interfaces at the CLI, on page 158.

For information about the Firewall Threat Defense CLI, see the Cisco Secure Firewall Threat Defense Command Reference.

Procedure

Step 1 If you are changing the data management interface to a new interface, move the current interface cable to the new interface.

Step 2 Connect to the device CLI.

You should use the console port when using these commands. If you are performing initial setup, then you may be disconnected from the Management interface. If you are editing the configuration due to a disrupted management connection, and you have SSH access to the dedicated Management interface, then you can use that SSH connection.

See Log Into the Command-Line Interface on the Device, on page 11.

Step 3 Log in with the **admin** username and password.

Step 4 Disable the interface so you can reconfigure its settings.

configure network management-data-interface disable

Note
If you only want to set a new IPv4 address on the same interface and not make any other changes, you can skip this step. Other changes require you to disable the interface first.

Example:

```
> configure network management-data-interface disable

 Configuration updated successfully..!!


Configuration disable was successful, please update the default route to point to a gateway
 on management interface using the command 'configure network'
```

Step 5 Configure the new data interface for manager access.

configure network management-data-interface

You are then prompted to configure basic network settings for the data interface.

If you change the data management interface to a new interface on the same network, use the same settings as for the previous interface except the interface ID. In addition, for the **Do you wish to clear all the device configuration before applying ? (y/n) [n]:** option, choose **y**. This choice will clear the old data management

interface configuration, so that you can successfully reuse the IP address and interface name on the new interface.

```
> configure network management-data-interface
Data interface to use for management: ethernet1/4
Specify a name for the interface [outside]: internet
IP address (manual / dhcp) [dhcp]: manual
IPv4/IPv6 address: 10.10.6.7
Netmask/IPv6 Prefix: 255.255.255.0
Default Gateway: 10.10.6.1
Comma-separated list of DNS servers [none]: 208.67.222.222,208.67.220.220
DDNS server update URL [none]:
Do you wish to clear all the device configuration before applying ? (y/n) [n]: y

Configuration done with option to allow manager access from any network, if you wish to
change the manager access network
use the 'client' option in the command 'configure network management-data-interface'.

Setting IPv4 network configuration.
Network settings changed.

>
```

Step 6 (Optional) Limit data interface access to the Firewall Management Center on a specific network.

configure network management-data-interface client *ip_address netmask*

By default, all networks are allowed.

Step 7 Update the Hostname or IP Address in the Firewall Management Center, on page 170.

The connection may be reestablished automatically, but disabling and reenabling the connection in the Firewall Management Center will help the connection reestablish faster. Or you may need to update the device IP address in the Firewall Management Center according to the linked procedure.

Step 8 Check that the management connection was reestablished.

sftunnel-status-brief

See the following sample output for a connection that is up, with peer channel and heartbeat information shown:

```
> sftunnel-status-brief
PEER:10.10.17.202
Peer channel Channel-A is valid type (CONTROL), using 'eth0', connected to '10.10.17.202'
via '10.10.17.222'
Peer channel Channel-B is valid type (EVENT), using 'eth0', connected to '10.10.17.202' via
 '10.10.17.222'
Registration: Completed.
IPv4 Connection to peer '10.10.17.202' Start Time: Wed Jun 10 14:27:12 2020 UTC
Heartbeat Send Time: Mon Jun 15 09:02:08 2020 UTC
Heartbeat Received Time: Mon Jun 15 09:02:16 2020 UTC
```

Step 9 In the Firewall Management Center, choose **Devices** > **Device Management**, and click **Edit** (). In the **Device** > **Management** area, click **Refresh** next to **Manager Access - Configuration Details**.

The Firewall Management Center detects the interface and default route configuration changes and blocks deployment to the device. When you change the data interface settings locally on the device, you must reconcile

those changes in the Firewall Management Center manually. You can view the discrepancies between the Firewall Management Center and the device on the **Configuration** tab.

Step 10 Choose **Interfaces**, and make the following changes.

a) Remove the IP address and name from the old data management interface and disable manager access for this interface.

b) Configure the new data management interface with the new settings (the ones you used at the CLI) and enable manager access for it.

Step 11 Choose **Routing** > **Static Route** and change the default route from the old data management interface to the new one.

Step 12 Return to the **Manager Access - Configuration Details** dialog box, and click **Acknowledge** to remove the deployment block.

The next time you deploy, the Firewall Management Center configuration will overwrite any remaining conflicting settings on the Firewall Threat Defense. It is your responsibility to manually fix the configuration in the Firewall Management Center before you re-deploy.

You will see expected messages of "Config was cleared" and "Manager access changed and acknowledged."

Modify the Firewall Threat Defense Data Interface Used for Management in the GUI

If the management connection is up, but you want to change the IP address of the data interface used for manager access, follow these steps. For example, if you register a device using zero-touch provisioning, then you need to change the IP address to a static address before you can enable high availability.

You can alternatively change interface settings at the CLI, but we recommend only using that method if the management connection is down. Any changes you make at the CLI will have to be replicated in the GUI anyway.

Procedure

Step 1 Choose **Devices** > **Device Management**, and click **Edit** (✎) next to the device.

Step 2 Choose **Interfaces**.

Step 3 If you want to change the interface used for manager access:

a) Remove the IP address and name from the old data management interface and disable manager access for this interface.

b) Configure the new data management interface with the new settings and enable manager access for it.

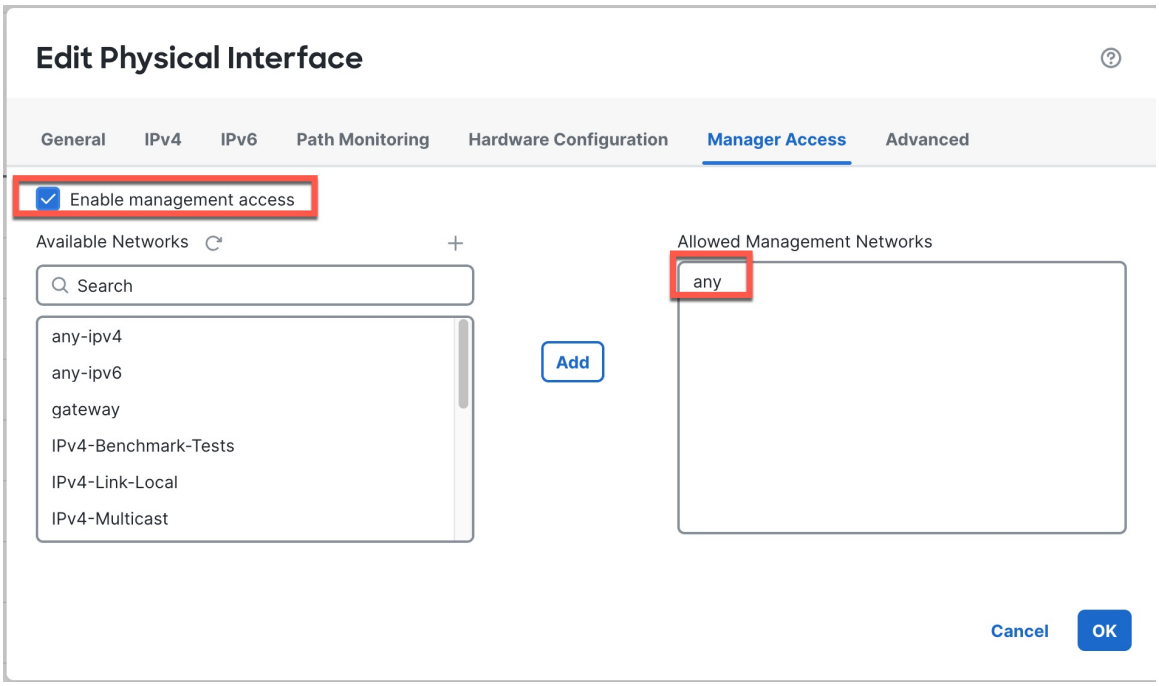

c) If you use a static IP address, you are reminded to make sure you have a default route. Click **Yes**.

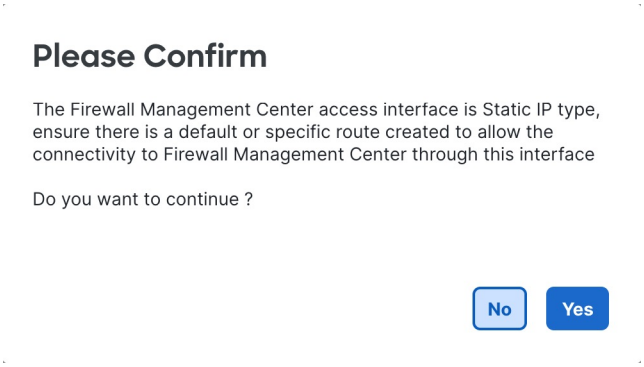

d) Click **OK** to exit the interface.

e) Click **Save** on the **Interfaces** page.

Step 4 If you only want to change the IP address:

a) Change the IP address.

b) For a static IP address, you are reminded to make sure you have a default route. Click **Yes**.

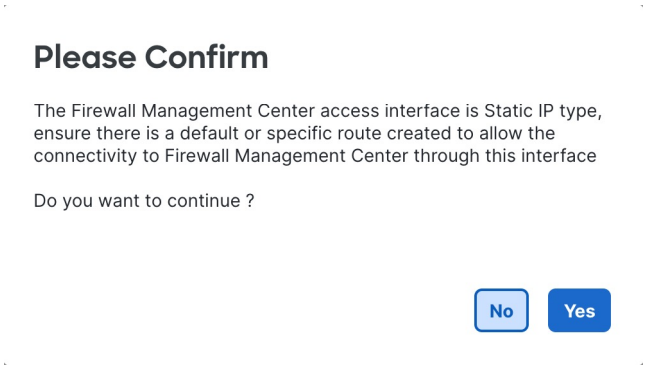

c) Click **OK** to exit the interface.
d) Click **Save** on the **Interfaces** page.

Step 5 Choose **Routing** > **Static Route** and add or change the default or static route for the manager access interface.

Step 6 Deploy configuration changes; see Deploy Configuration Changes, on page 259.

The Firewall Management Center will deploy the configuration changes over the current connection. After the deployment, the data interface will have a new IP address, so the management connection will need to be reestablished.

Step 7 Update the Hostname or IP Address in the Firewall Management Center, on page 170.

Step 8 Ensure the management connection is reestablished.

In the **Device** > **Management** area, click **Manager Access Details: Configuration** and then click **Connection Status**.

The following status shows a successful connection for a data interface, showing the internal "tap_nlp" interface.

Figure 68: Connection Status

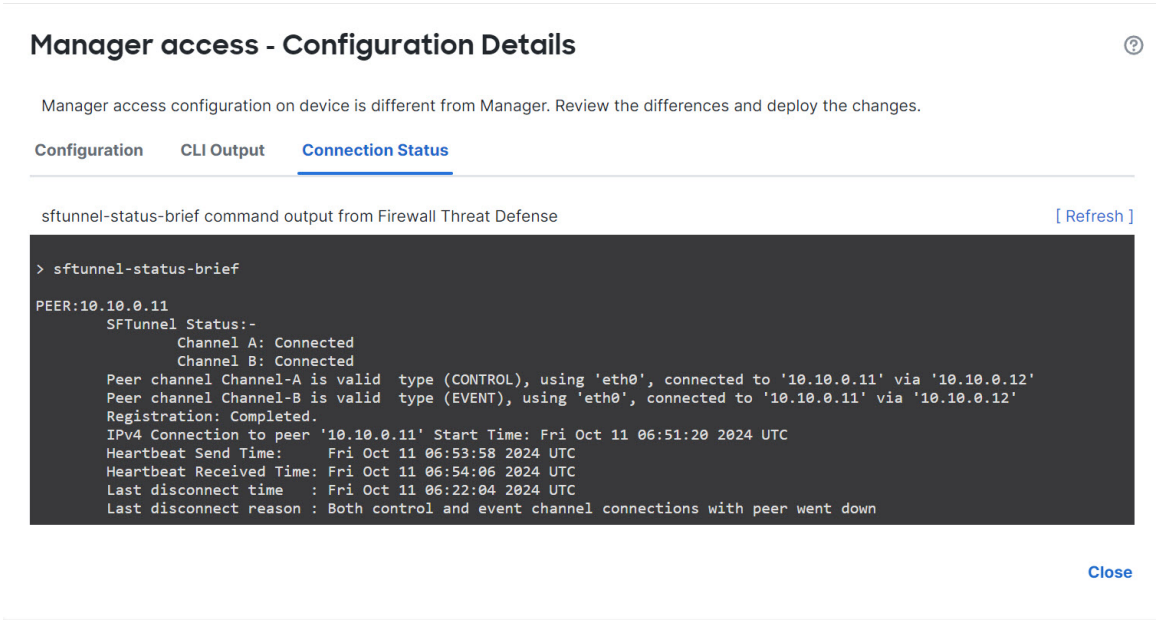

If it takes more than 10 minutes to reestablish the connection, you should troubleshoot the connection. See Troubleshoot Management Connectivity on a Data Interface, on page 190.

Update the Hostname or IP Address in the Firewall Management Center

If you edit the hostname or IP address of a device after you added it to the Firewall Management Center (using the device's CLI, for example), you need to use the procedure below to manually update the hostname or IP address on the managing Firewall Management Center.

To change the device management IP address on the device, see Modify Firewall Threat Defense Management Interfaces at the CLI, on page 158.

If you used only the NAT ID when registering the device, then the IP shows as **NO-IP** on this page, and you do not need to update the IP address/hostname.

If you used zero-touch provisioning to register the device on the outside interface, the hostname is automatically generated along with a matching DDNS configuration; you cannot edit the hostname in this case.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to modify management options, click **Edit** ( ).

Step 3 Click **Device**, and view the **Management** area.

Step 4 Disable management temporarily by clicking the slider so it is disabled **Slider disabled** ( ).

Figure 69: Disable Management

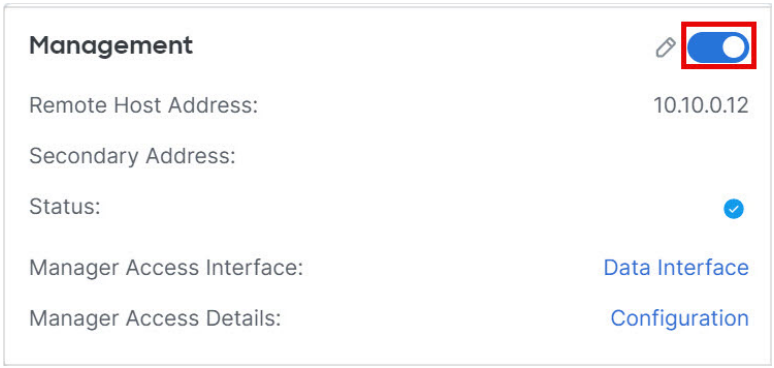

You are prompted to proceed with disabling management; click **Yes**.

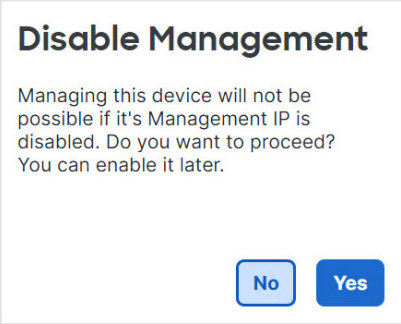

Disabling management blocks the connection between the Firewall Management Center and the device, but does **not** unregister the device from the Firewall Management Center.

Step 5 Edit the **Remote Host Address** IP address and optional **Secondary Address** (when using a redundant data interface) or hostname by clicking **Edit** ( ).

Figure 70: Edit Management Address

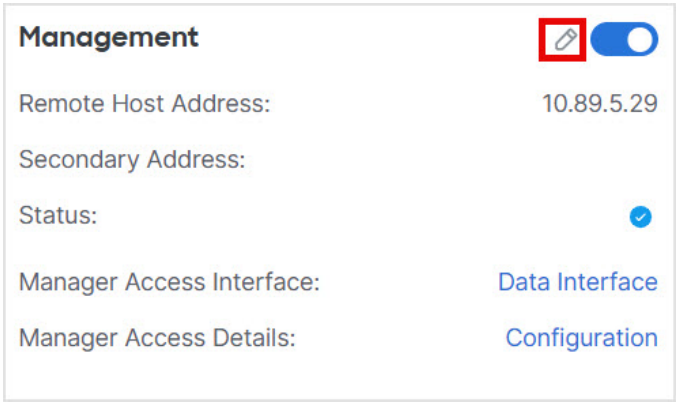

Step 6 In the **Management** dialog box, modify the name or IP address in the **Remote Host Address** field and the optional **Secondary Address** field, and click **Save**.

For information about using a secondary manager access data interface, see Configure a Redundant Manager Access Data Interface, on page 153.

Figure 71: Management IP Address

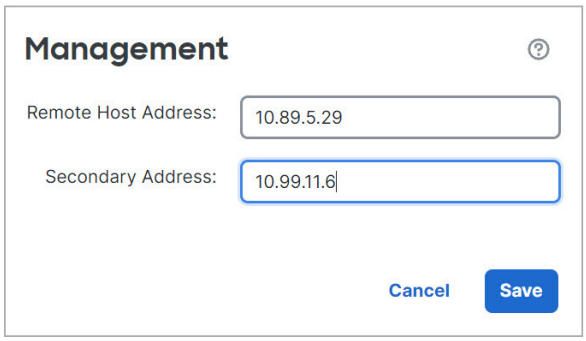

Step 7 Reenable management by clicking the slider so it is enabled **Slider enabled** ( ).

Figure 72: Enable Management Connection

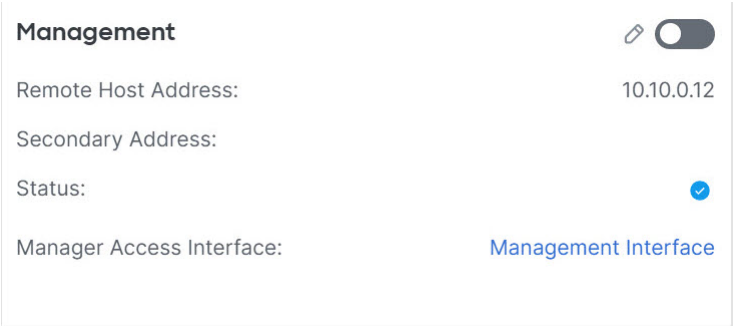

Change the Firewall Management Center IP Address

If you change the Firewall Management Center IP address or hostname, you should also change the value at the device CLI so the configurations match. Although in most cases, the management connection will be reestablished without changing the Firewall Management Center IP address or hostname on the device, in at least one case, you must perform this task for the connection to be reestablished: when you added the device to the Firewall Management Center and you specified the NAT ID only. Even in other cases, we recommend keeping the Firewall Management Center IP address or hostname up to date for extra network resiliency.

Procedure

Step 1 Change the Firewall Management Center IP address.

Caution
Be careful when making changes to the Firewall Management Center interface to which you are connected; if you cannot re-connect because of a configuration error, you need to access the Firewall Management Center console port to re-configure the network settings in the Linux shell. You must contact Cisco TAC to guide you in this operation.

a) Choose **System** (⚙) > **Configuration** > **Management Interfaces**.
b) In the **Interfaces** area, click **Edit** next to the interface that you want to configure.
c) Change the IP address, and click **Save**.

Step 2 At the Firewall Threat Defense CLI, view the Firewall Management Center identifier.

show managers

Example:

```
> show managers
Type                      : Manager
Host                      : 10.10.1.4
Display name              : 10.10.1.4
Identifier                : f7ffad78-bf16-11ec-a737-baa2f76ef602
Registration              : Completed
Management type           : Configuration
```

Step 3 At the Firewall Threat Defense CLI, edit the Firewall Management Center IP address or hostname.

configure manager edit *identifier* {**hostname** {*ip_address* | *hostname*} | **displayname** *display_name*}

If the Firewall Management Center was originally identified by **DONTRESOLVE** and a NAT ID, you can change the value to a hostname or IP address using this command. You cannot change an IP address or hostname to **DONTRESOLVE**.

The management connection will go down, and then reestablish. You can monitor the state of the connection using the **sftunnel-status** command.

Example:

```
> configure manager edit f7ffad78-bf16-11ec-a737-baa2f76ef602 hostname 10.10.5.1
```

Change Both Firewall Management Center and Threat Defense IP Addresses

You might want to change both Firewall Management Center and Firewall Threat Defense IP addresses if you need to move them to a new network.

Procedure

Step 1 Disable the management connection.

For a high-availability pair or cluster, perform these steps on all units.

a) Choose **Devices** > **Device Management**.
b) Next to the device, click **Edit** ( ).
c) Click **Device**, and view the **Management** area.
d) Disable management temporarily by clicking the slider so it is disabled ( ).

Figure 73: Disable Management

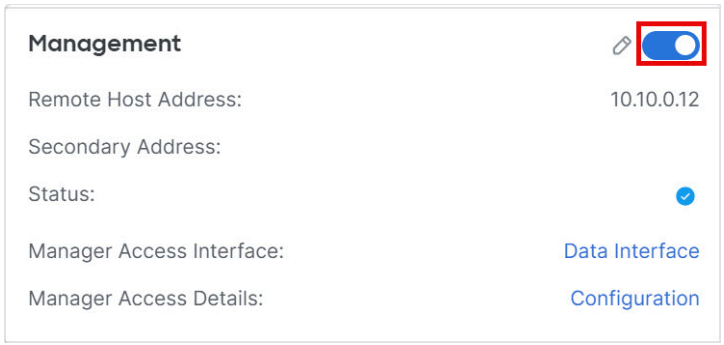

You are prompted to proceed with disabling management; click **Yes**.

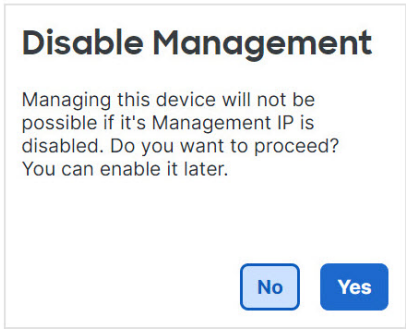

Step 2 Change the device IP address in the Firewall Management Center to the new device IP address.

You will change the IP address on the device later.

For a high-availability pair or cluster, perform these steps on all units.

a) Edit the **Remote Host Address** IP address and optional **Secondary Address** (when using a redundant data interface) or hostname by clicking **Edit** ( ).

Figure 74: Edit Management Address

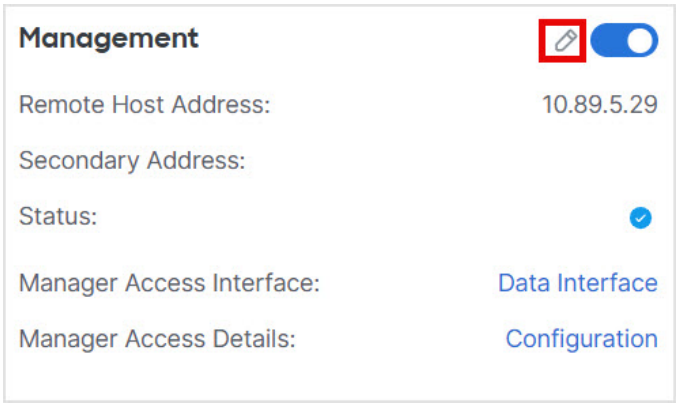

b) In the **Management** dialog box, modify the name or IP address in the **Remote Host Address** field and the optional **Secondary Address** field, and click **Save**.

Figure 75: Management IP Address

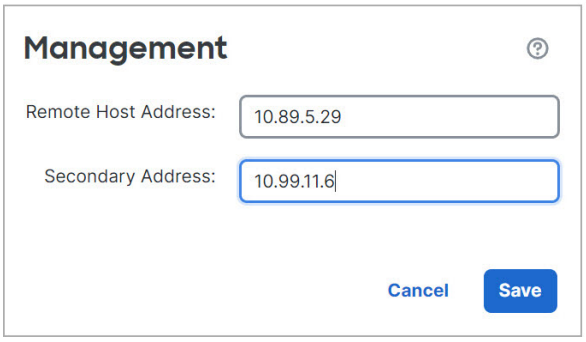

Step 3 Change the Firewall Management Center IP address.

Caution
Be careful when making changes to the Firewall Management Center interface to which you are connected; if you cannot re-connect because of a configuration error, you need to access the Firewall Management Center console port to re-configure the network settings in the Linux shell. You must contact Cisco TAC to guide you in this operation.

a) Choose **System** (⚙) > **Configuration** > **Management Interfaces**.
b) In the **Interfaces** area, click **Edit** next to the interface that you want to configure.
c) Change the IP address, and click **Save**.

Step 4 Change the manager IP address on the device.

For a high-availability pair or cluster, perform these steps on all units.

a) At the Firewall Threat Defense CLI, view the Firewall Management Center identifier.

show managers

Example:

```
> show managers
Type                      : Manager
Host                      : 10.10.1.4
Display name              : 10.10.1.4
Identifier                : f7ffad78-bf16-11ec-a737-baa2f76ef602
Registration              : Completed
Management type           : Configuration
```

b) Edit the Firewall Management Center IP address or hostname.

configure manager edit *identifier* {**hostname** {*ip_address* | *hostname*} | **displayname** *display_name*}

If the Firewall Management Center was originally identified by **DONTRESOLVE** and a NAT ID, you can change the value to a hostname or IP address using this command. You cannot change an IP address or hostname to **DONTRESOLVE**.

Example:

```
> configure manager edit f7ffad78-bf16-11ec-a737-baa2f76ef602 hostname 10.10.5.1
```

Step 5 Change the IP address of the manager access interface at the console port.

For a high-availability pair or cluster, perform these steps on all units.

If you use the dedicated Management interface:

configure network ipv4

configure network ipv6

If you use the dedicated Management interface:

configure network management-data-interface disable

configure network management-data-interface

Step 6 Reenable management by clicking the slider so it is enabled (🔵).

For a high-availability pair or cluster, perform these steps on all units.

Figure 76: Enable Management Connection

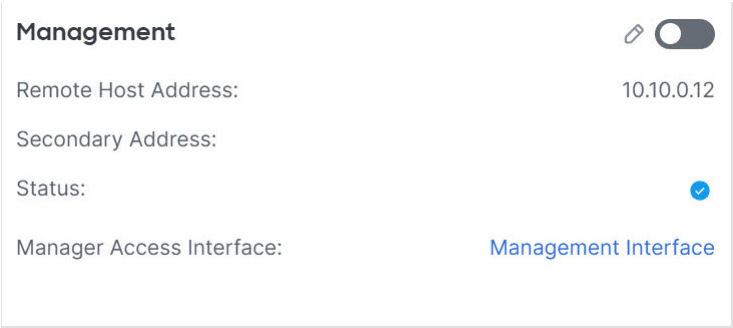

Step 7 (If using a data interface for manager access) Refresh the data interface settings in the Firewall Management Center.

For a high-availability pair, perform this step on both units.

a) Choose **Devices** > **Device Management**, and click **Manager Access - Configuration Details**, and then click **Refresh**.
b) Choose **Devices** > **Device Management**, and click the **Interfaces** tab and set the IP address to match the new address.
c) Return to the **Manager Access - Configuration Details** dialog box, and click **Acknowledge** to remove the deployment block.

Step 8 Ensure the management connection is reestablished.

In the Firewall Management Center, check the management connection status. Navigate to the **Devices** > **Device Management**, and click **Management** section under the **Device** tab. Then, click **Manager Access - Configuration Details** to view the **Connection Status** page.

At the Firewall Threat Defense CLI, enter the **sftunnel-status-brief** command to view the management connection status.

The following status shows a successful connection for a data interface, showing the internal "tap_nlp" interface.

Figure 77: Connection Status

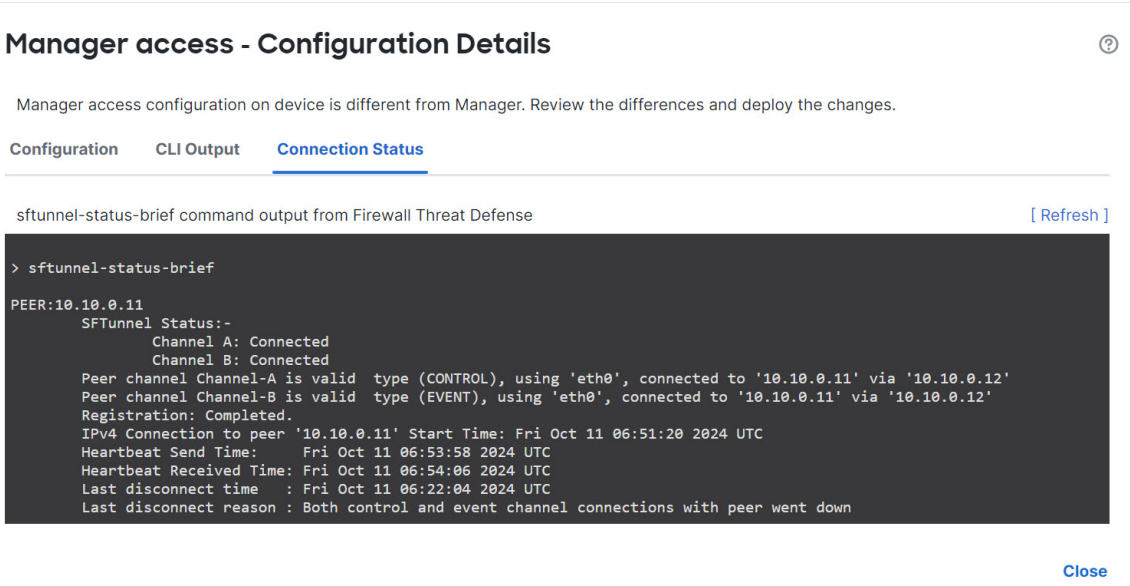

Step 9 (For a high-availability Firewall Management Center pair) Repeat configuration changes on the secondary Firewall Management Center.

a) Change the secondary Firewall Management Center IP address.
b) Specify the new peer addresses on both units.
c) Make the secondary unit the active unit.
d) Disable the device management connection.
e) Change the device IP address in the Firewall Management Center.
f) Reenable the management connection.

Change the Manager Access Interface

After you register the device, you can change the manager access interface, between the Management interface a data interface.

Change the Manager Access Interface from Management to Data

You can manage the Firewall Threat Defense from either the dedicated Management interface or from a data interface. If you want to change the manager access interface after you added the device to the Firewall Management Center, follow these steps to migrate from the Management interface to a data interface. To migrate the other direction, see Change the Manager Access Interface from Data to Management, on page 182.

Initiating the manager access migration from Management to data causes the Firewall Management Center to apply a block on deployment to the Firewall Threat Defense. To remove the block, enable manager access on the data interface.

See the following steps to enable manager access on a data interface and also configure other required settings.

Before you begin

For high-availability pairs, unless stated otherwise, perform all steps only on the active unit. Once the configuration changes are deployed, the standby unit synchronizes configuration and other state information from the active unit.

Procedure

Step 1 Initiate the interface migration.

a) On the **Devices** > **Device Management** page, click **Edit** (✎) for the device. Click **Device**, and in the **Management** area, click the **Manager Access Interface** link.

The **Manager Access Interface** field shows the current Management interface. When you click the link, choose the new interface type, **Data Interface**, in the **Manage device by** drop-down list.

Figure 78: Manager Access Interface

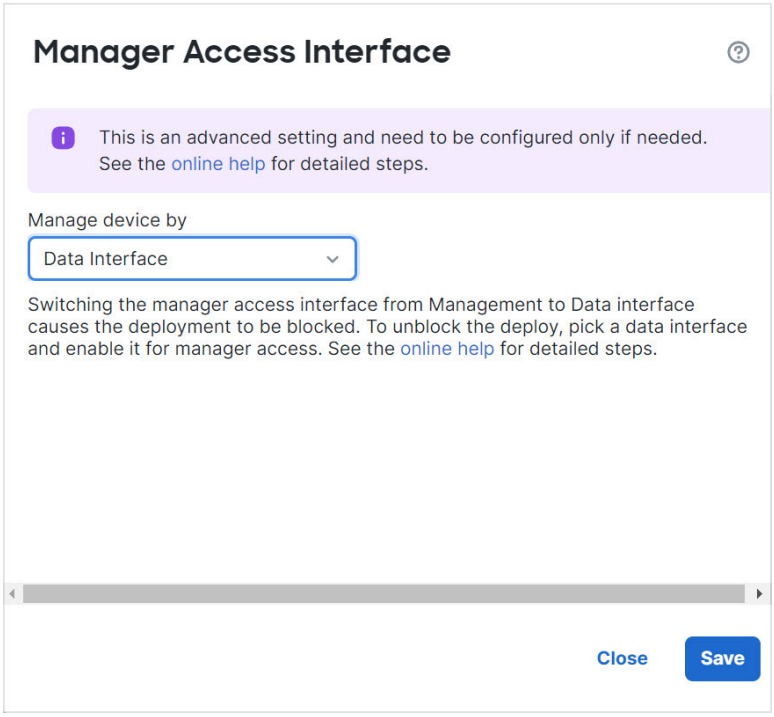

b) Click **OK** and then **Close**.

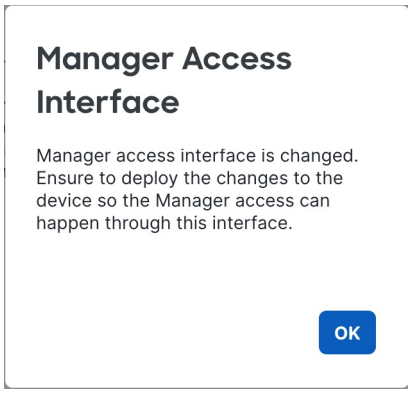

You must now complete the remaining steps in this procedure to enable manager access on the data interface. The **Management** area now shows **Manager Access Interface: Data Interface**, and **Manager Access Details: Configuration**.

Figure 79: Manager Access

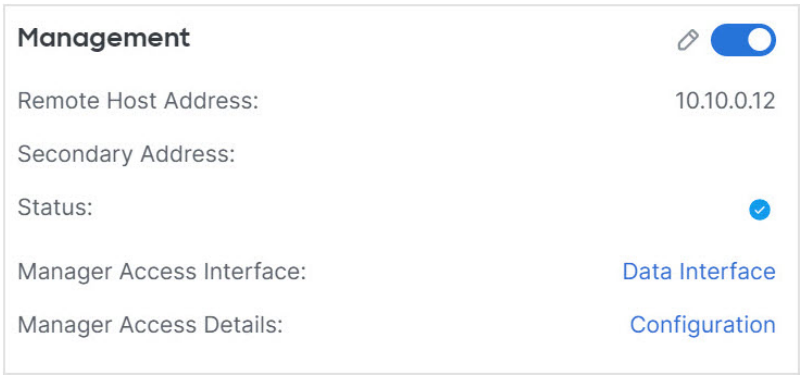

If you click **Configuration**, the **Manager Access - Configuration Details** dialog box opens. The **Manager Access Mode** shows a Deploy pending state.

Step 2 Enable manager access on the data interface(s). Click **Interfaces**, click **Edit** (✎) for the interface, and then click **Manager Access**.

Check **Enable management access** and click **OK**. By default, all networks are allowed, but you can limit access as long as the Firewall Management Center address is allowed.

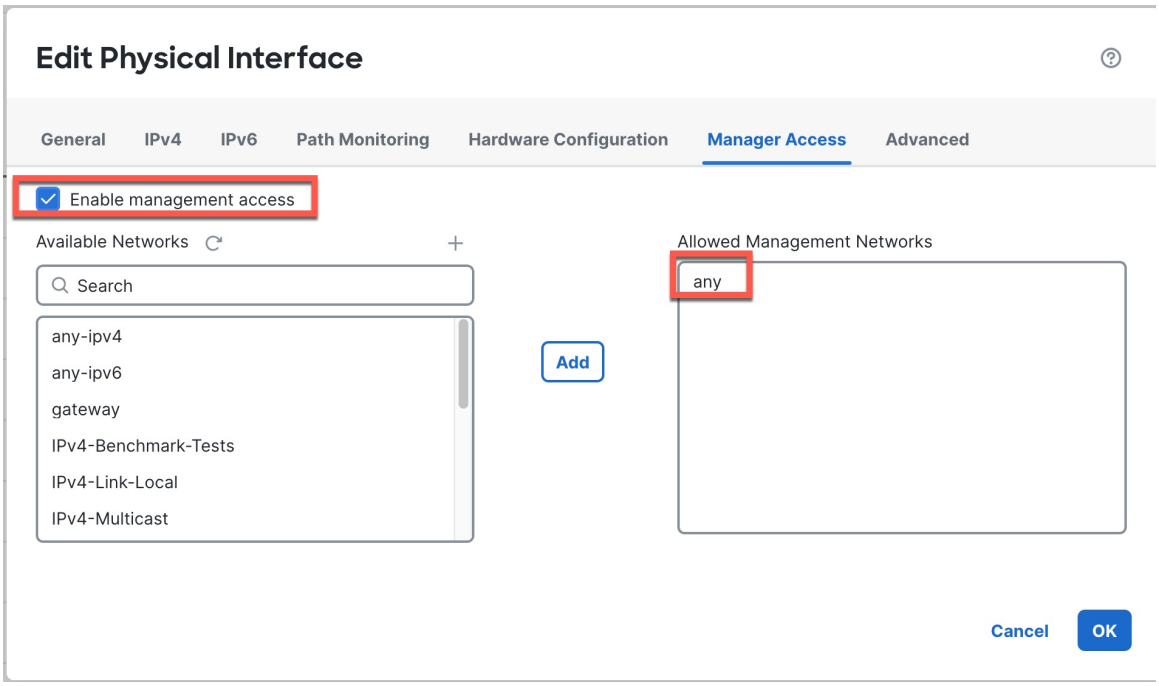

If the manager access interface uses a static IP address, you are reminded to configure routing for it.

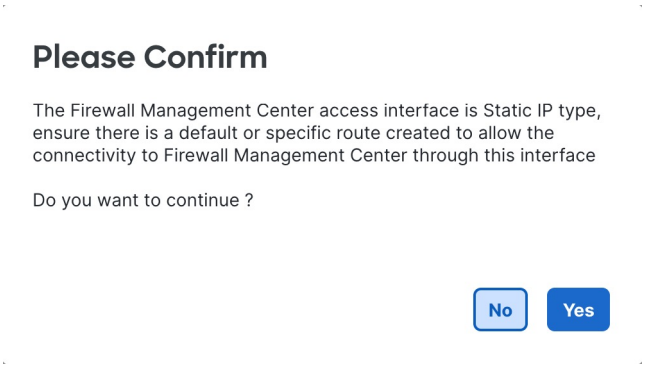

Click **Save** on the **Interfaces** page. See Configure Routed Mode Interfaces, on page 856 for more information about interface settings. You can enable manager access on one routed data interface, plus an optional secondary interface. Make sure these interfaces are fully configured with a name and IP address and that they are enabled.

If you use a secondary interface for redundancy, see Configure a Redundant Manager Access Data Interface, on page 153 for additional required configuration.

Step 3 (Optional) If you use DHCP for the interface, enable the web type DDNS method on the **DDNS** page. Navigate to **Devices** > **Device Management**, and then click **DDNS** under the **DHCP** tab.

See Configure Dynamic DNS, on page 920. DDNS ensures the Firewall Management Center can reach the Firewall Threat Defense at its Fully-Qualified Domain Name (FQDN) if the FTD's IP address changes.

Step 4 Make sure the Firewall Threat Defense can route to the Firewall Management Center through the data interface; add a static route if necessary on the Static Route page. Navigate to **Devices** > **Device Management** and then click **Static Route** under the **Routing** tab.

See Add a Static Route, on page 1154.

Step 5 (Optional) Configure DNS in a Platform Settings policy, and apply it to this device at **Devices** > **Platform Settings** > **DNS**.

See DNS, on page 949. DNS is required if you use DDNS. You may also use DNS for FQDNs in your security policies.

Step 6 (Optional) Enable SSH for the data interface in a Platform Settings policy, and apply it to this device at **Devices** > **Device Management** page. Click **Edit** (✎) for the device and then click **SSH Access**.

See SSH Access, on page 964. SSH is not enabled by default on the data interfaces, so if you want to manage the Firewall Threat Defense using SSH, you need to explicitly allow it.

Step 7 Deploy configuration changes; see Deploy Configuration Changes, on page 259.

You will see a validation error to confirm that you are changing the manager access interface. Check **Ignore warnings** and deploy again.

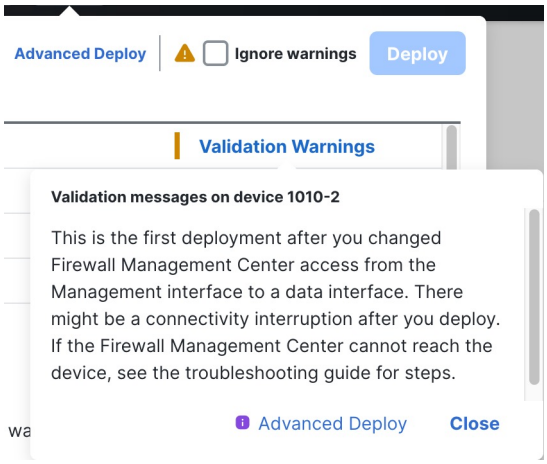

The Firewall Management Center will deploy the configuration changes over the current Management interface. After the deployment, the data interface is now ready for use, but the original management connection to Management is still active.

Step 8 At the Firewall Threat Defense CLI (preferably from the console port), set the Management interface to use a static IP address and set the gateway to use the data interfaces. For high availability, perform this step on both units.

configure network {**ipv4** | **ipv6**} **manual** *ip_address netmask* **data-interfaces**

- *ip_address netmask*—Although you do not plan to use the Management interface, you must set a static IP address, for example, a private address so that you can set the gateway to **data-interfaces** (see the next bullet). You cannot use DHCP because the default route, which must be **data-interfaces**, might be overwritten with one received from the DHCP server.
- **data-interfaces**—This setting forwards management traffic over the backplane so it can be routed through the manager access data interface.

We recommend that you use the console port instead of an SSH connection because when you change the Management interface network settings, your SSH session will be disconnected.

Step 9 If necessary, re-cable the Firewall Threat Defense so it can reach the Firewall Management Center on the data interface. For high availability, perform this step on both units.

Step 10 In the Firewall Management Center, disable the management connection, update the **Remote Host Address** IP address and optional **Secondary Address** for the Firewall Threat Defense in the **Devices** > **Device Management** page in the **Management** area under the **Device** tab, and reenable the connection.

See Update the Hostname or IP Address in the Firewall Management Center, on page 170. If you used the Firewall Threat Defense hostname or just the NAT ID when you added the Firewall Threat Defense to the Firewall Management Center, you do not need to update the value; however, you need to disable and reenable the management connection to restart the connection.

Step 11 Ensure the management connection is reestablished.

In the **Devices** > **Device Management** page, click **Manager Access Details: Configuration** and then click **Connection Status**.

Alternatively, you can check at the Firewall Threat Defense CLI. Enter the **sftunnel-status-brief** command to view the management connection status.

The following status shows a successful connection for a data interface, showing the internal "tap_nlp" interface.

Figure 80: Connection Status

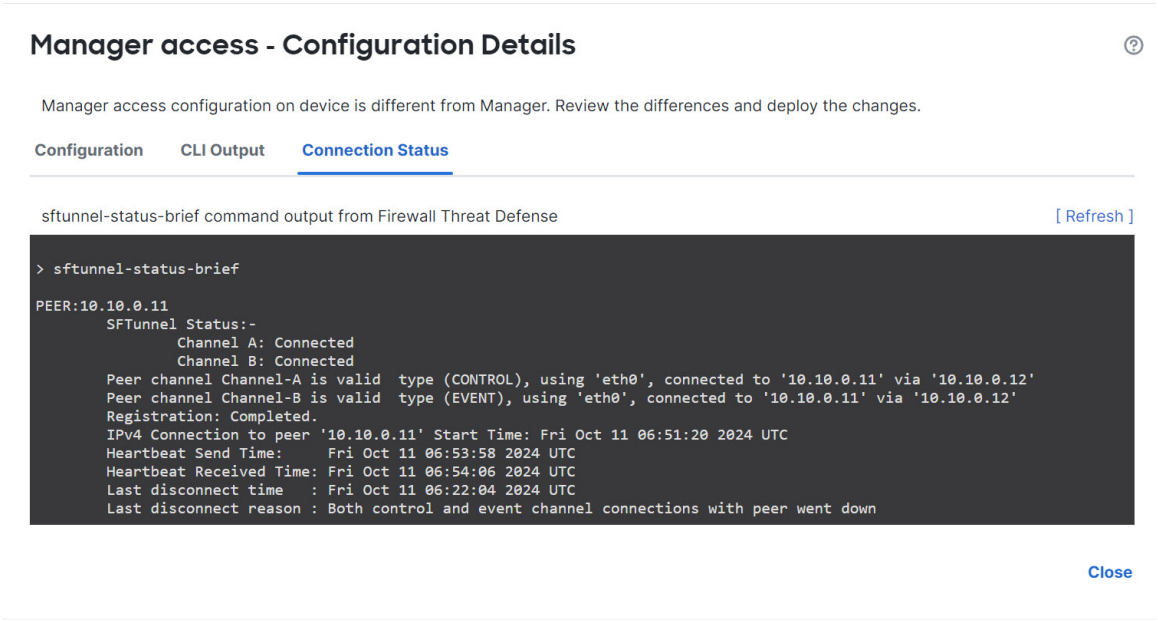

If it takes more than 10 minutes to reestablish the connection, you should troubleshoot the connection. See Troubleshoot Management Connectivity on a Data Interface, on page 190.

Change the Manager Access Interface from Data to Management

You can manage the Firewall Threat Defense from either the dedicated Management interface or from a data interface. If you want to change the manager access interface after you added the device to the Firewall Management Center, follow these steps to migrate from a data interface to the Management interface. To migrate the other direction, see Change the Manager Access Interface from Management to Data, on page 177.

Initiating the manager access migration from data to Management causes the Firewall Management Center to apply a block on deployment to the Firewall Threat Defense. You must disable manager access on the data interface to remove the block.

See the following steps to disable manager access on a data interface, and also configure other required settings.

Before you begin

For high-availability pairs, unless stated otherwise, perform all steps only on the active unit. Once the configuration changes are deployed, the standby unit synchronizes configuration and other state information from the active unit.

Procedure

Step 1 Initiate the interface migration.

a) On the **Devices** > **Device Management** page, click **Edit** (✎) for the device. Click **Device**, and in the **Management** area, click the link for **Manager Access Interface**.

The **Manager Access Interface** field shows the current management interface as data. When you click the link, choose the new interface type, **Management Interface**, in the **Manage device by** drop-down list.

Figure 81: Manager Access Interface

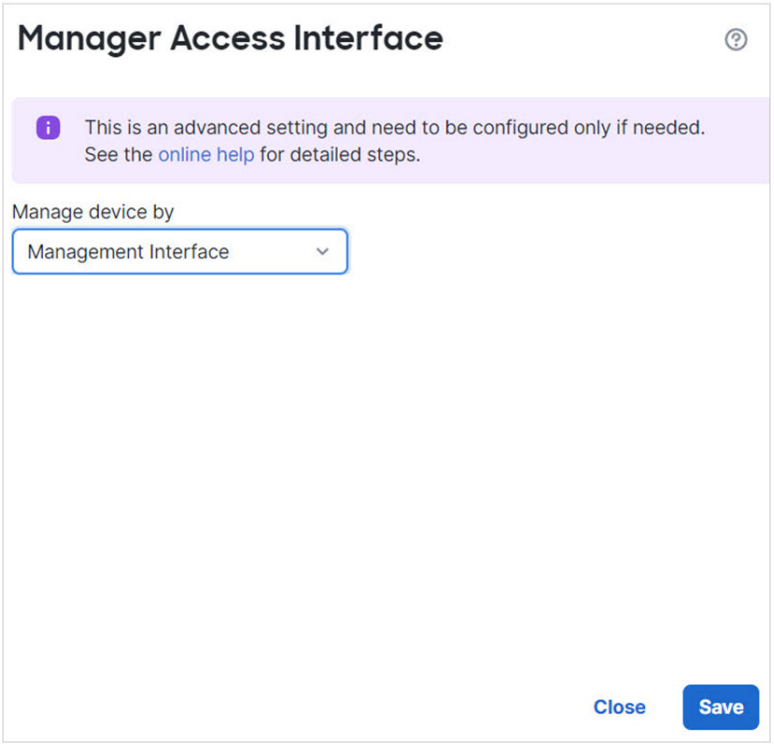

b) Click **Save**.

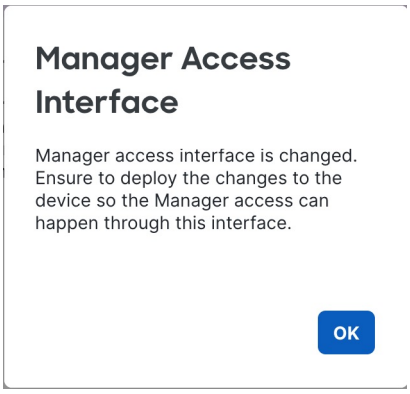

Click **OK** and then **Close**.

You must now complete the remaining steps in this procedure to enable manager access on the Management interface. The **Management** area now shows the **Manager Access Interface: Management Interface**.

Figure 82: Manager Access

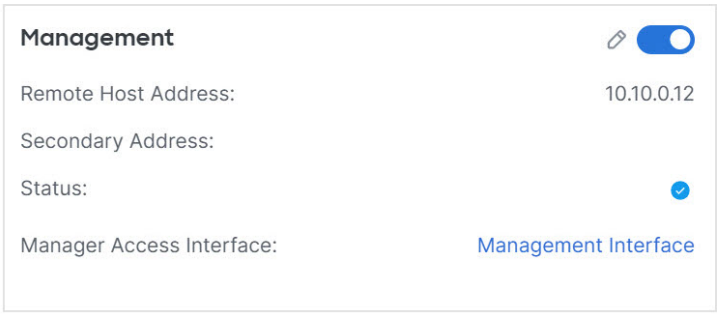

Step 2 Disable manager access on the data interface(s). Click **Interfaces**, click **Edit** (✎) for the interface, and then click **Manager Access**.

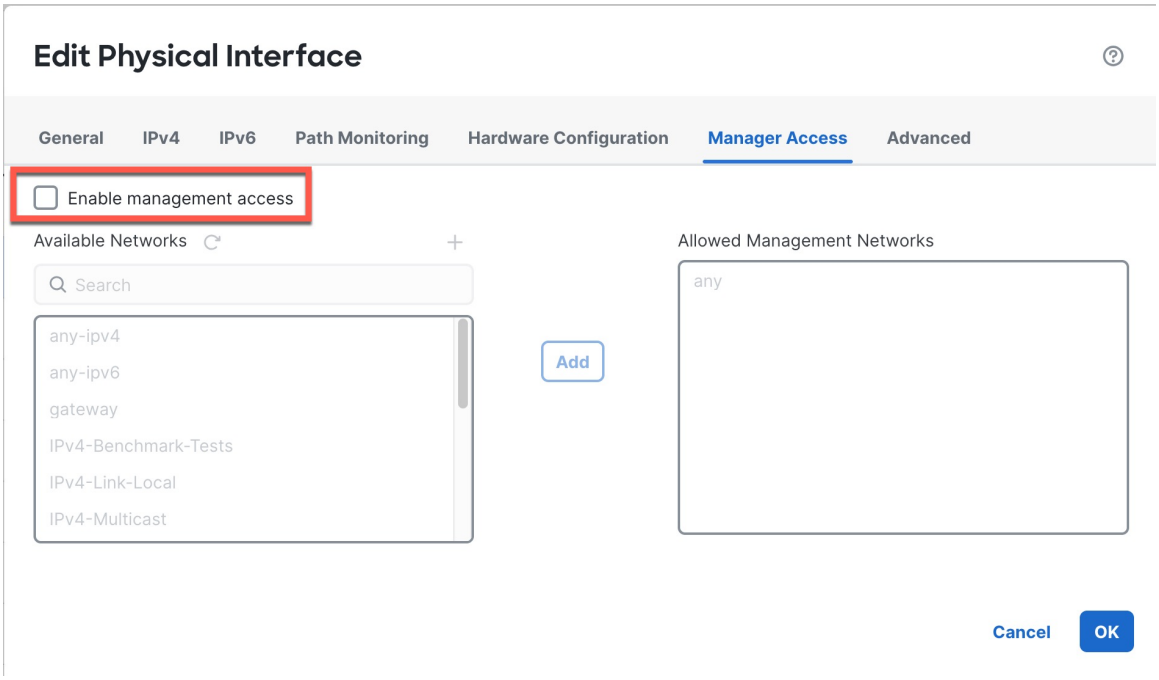

Uncheck **Enable management access** and click **OK**. Click **Save** on the **Interfaces** page. This step removes the block on deployment.

Step 3 If you have not already done so, configure DNS settings for the data interface in a Platform Setting policy, and apply it to this device at **Devices** > **Device Management** page. Click **Edit** (✎) for the device and then click **DNS**.

See DNS, on page 949. The Firewall Management Center deployment that disables manager access on the data interface will remove any local DNS configuration. If that DNS server is used in any security policy, such as an FQDN in an Access Rule, then you must re-apply the DNS configuration using the Firewall Management Center.

Step 4 Deploy configuration changes; see Deploy Configuration Changes, on page 259.

The Firewall Management Center will deploy the configuration changes over the current data interface.

Step 5 If necessary, re-cable the Firewall Threat Defense so it can reach the Firewall Management Center on the Management interface. For High Availability, perform this step on both units.

Step 6 At the Firewall Threat Defense CLI, configure the Management interface IP address and gateway using a static IP address or DHCP. For high availability, perform this step on both units.

When you originally configured the data interface for manager access, the Management gateway was set to data-interfaces, which forwarded management traffic over the backplane so it could be routed through the manager access data interface. You now need to set an IP address for the gateway on the management network.

Static IP address:

configure network {**ipv4** | **ipv6**} **manual** *ip_address netmask gateway_ip*

DHCP:

configure network{**ipv4** | **ipv6**} **dhcp**

Step 7 In the Firewall Management Center, disable the management connection, update the **Remote Host Address** IP address and optional **Secondary Address** for the Firewall Threat Defense in the **Devices** > **Device Management**, **Management** section under **Device** tab, and reenable the connection.

See Update the Hostname or IP Address in the Firewall Management Center, on page 170. If you used the Firewall Threat Defense hostname or just the NAT ID when you added the Firewall Threat Defense to the Firewall Management Center, you do not need to update the value; however, you need to disable and reenable the management connection to restart the connection.

Step 8 Ensure the management connection is reestablished.

In the Firewall Management Center, check the management connection status on the **Devices** > **Device Management**, **Management** section under **Device** tab, **Status** field or view notifications in the Firewall Management Center.

At the Firewall Threat Defense CLI, enter the **sftunnel-status-brief** command to view the management connection status.

If it takes more than 10 minutes to reestablish the connection, you should troubleshoot the connection. See Troubleshoot Management Connectivity on a Data Interface, on page 190.

View Manager Access Details for Data Interface Management

When you use a data interface for Firewall Management Center management instead of using the dedicated Management interface, you must be careful about changing the interface and network settings for the device in the Firewall Management Center so you do not disrupt the connection. You can also change the data interface settings locally on the device, which requires you to reconcile those changes in the Firewall Management Center manually. The **Devices** > **Device Management** > **Device** > **Management** > **Manager Access - Configuration Details** dialog box helps you resolve any discrepancies between the Firewall Management Center and the Firewall Threat Defense local configuration.

Normally, you configure the manager access data interface as part of initial Firewall Threat Defense setup before you add the Firewall Threat Defense to the Firewall Management Center. When you add the Firewall Threat Defense to the Firewall Management Center, the Firewall Management Center discovers and maintains the interface configuration, including the following settings: interface name and IP address, static route to the gateway, DNS servers, and DDNS server. For the DNS server, the configuration is maintained locally if it is discovered during registration, but it is not added to the Platform Settings policy in Firewall Management Center.

After you add the Firewall Threat Defense to the Firewall Management Center, if you change the data interface settings on the Firewall Threat Defense locally using the **configure network management-data-interface** command, then the Firewall Management Center detects the configuration changes, and blocks deployment to the Firewall Threat Defense. The Firewall Management Center detects the configuration changes using one of the following methods:

- Deploy to the Firewall Threat Defense. Before the Firewall Management Center deploys, it will detect the configuration differences and stop the deployment.
- The **Refresh** button on the **Manager Access - Configuration Details** dialog box.

To remove the block, you must go to the **Manager Access - Configuration Details** dialog box and click **Acknowledge**. The next time you deploy, the Firewall Management Center configuration will overwrite any remaining conflicting settings on the Firewall Threat Defense. It is your responsibility to manually fix the configuration in the Firewall Management Center before you re-deploy.

See the following pages on this dialog box.

Configuration

View the configuration comparison of the manager access data interface on the Firewall Management Center and the Firewall Threat Defense.

The following example shows the configuration details of the Firewall Threat Defense where the **configure network management-data-interface** command was entered on the Firewall Threat Defense. The pink highlights show that if you **Acknowledge** the differences but do not match the configuration in the Firewall Management Center, then the Firewall Threat Defense configuration will be removed. The blue highlights show configurations that will be modified on the Firewall Threat Defense. The green highlights show configurations that will be added to the Firewall Threat Defense.

The following example shows this page after configuring the interface in the Firewall Management Center; the interface settings match, and the pink highlight was removed.

CLI Output

View the CLI configuration of the manager access data interface, which is useful if you are familiar with the underlying CLI.

Figure 83: CLI Output

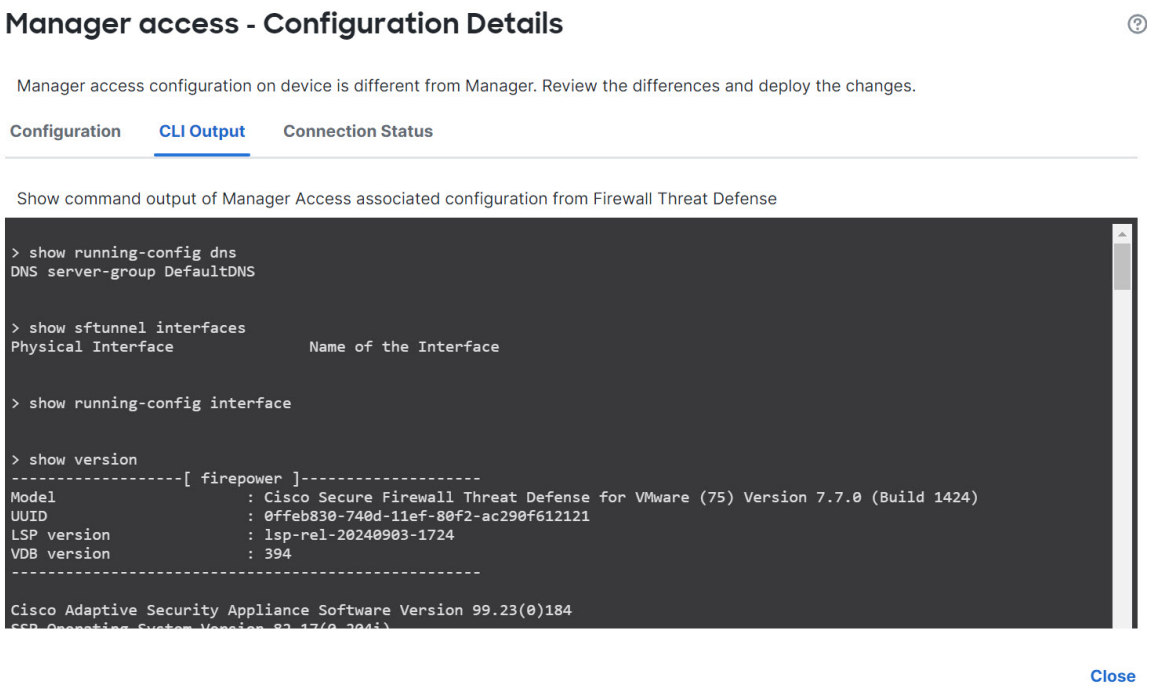

Connection Status

View management connection status. The following example shows that the management connection is still using the Management "management0" interface.

Figure 84: Connection Status

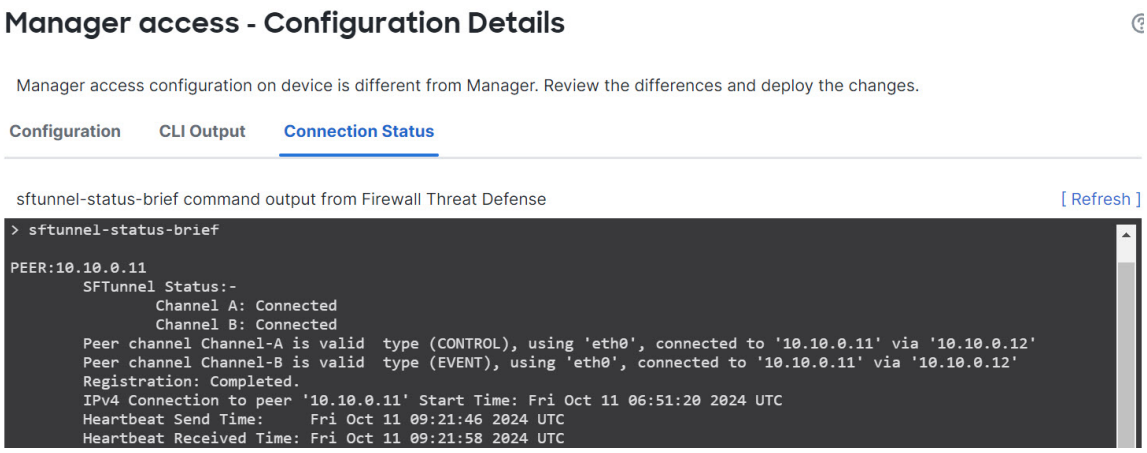

The following status shows a successful connection for a data interface, showing the internal "tap_nlp" interface.

Figure 85: Connection Status

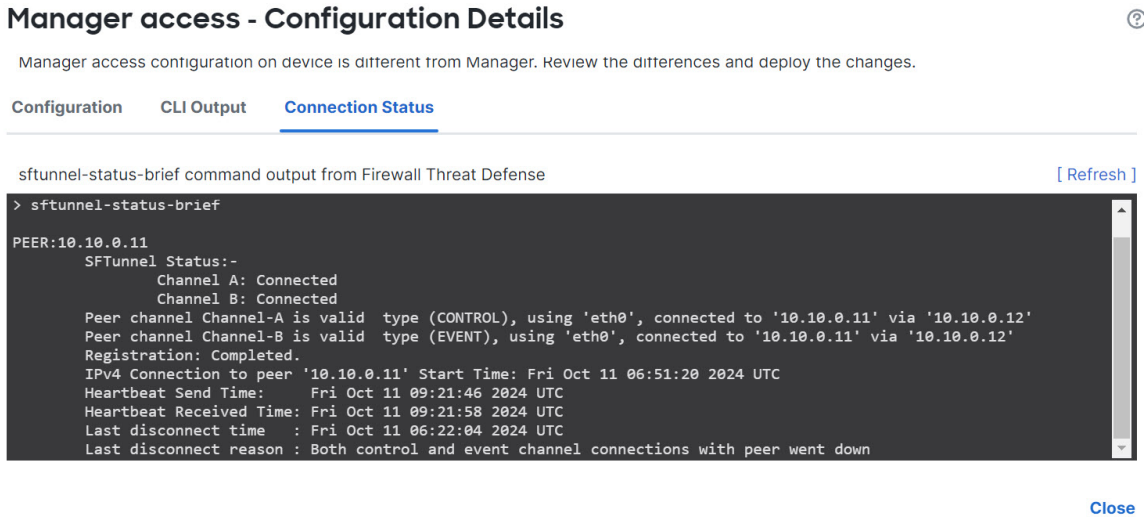

See the following sample output for a connection that is down; there is no peer channel "connected to" information, nor heartbeat information shown:

```
> sftunnel-status-brief
PEER:10.10.17.202
Registration: Completed.
Connection to peer '10.10.17.202' Attempted at Mon Jun 15 09:21:57 2020 UTC
Last disconnect time : Mon Jun 15 09:19:09 2020 UTC
Last disconnect reason : Both control and event channel connections with peer went down
```

See the following sample output for a connection that is up, with peer channel and heartbeat information shown:

```
> sftunnel-status-brief
PEER:10.10.17.202
Peer channel Channel-A is valid type (CONTROL), using 'eth0', connected to '10.10.17.202'
via '10.10.17.222'
Peer channel Channel-B is valid type (EVENT), using 'eth0', connected to '10.10.17.202' via
 '10.10.17.222'
Registration: Completed.
IPv4 Connection to peer '10.10.17.202' Start Time: Wed Jun 10 14:27:12 2020 UTC
Heartbeat Send Time: Mon Jun 15 09:02:08 2020 UTC
Heartbeat Received Time: Mon Jun 15 09:02:16 2020 UTC
```

Troubleshooting the Management Connection

Manually Roll Back the Configuration if the Firewall Management Center Loses Connectivity

If you use a data interface on the Firewall Threat Defense for manager access, and you deploy a configuration change from the Firewall Management Center that affects the network connectivity, you can roll back the configuration on the Firewall Threat Defense to the last-deployed configuration so you can restore management connectivity. You can then adjust the configuration settings in Firewall Management Center so that the network connectivity is maintained, and re-deploy. You can use the rollback feature even if you do not lose connectivity; it is not limited to this troubleshooting situation.

Alternatively, you can enable auto rollback of the configuration if you lose connectivity after a deployment; see Edit Deployment Settings, on page 201.

See the following guidelines:

- Only the previous deployment is available locally on the Firewall Threat Defense; you cannot roll back to any earlier deployments.
- Rollback is supported for high availability but not supported for clustering deployments.
- The rollback only affects configurations that you can set in the Firewall Management Center. For example, the rollback does not affect any local configuration related to the dedicated Management interface, which you can only configure at the Firewall Threat Defense CLI. Note that if you changed data interface settings after the last Firewall Management Center deployment using the **configure network management-data-interface** command, and then you use the rollback command, those settings will not be preserved; they will roll back to the last-deployed Firewall Management Center settings.
- UCAPL/CC mode cannot be rolled back.
- Out-of-band SCEP certificate data that was updated during the previous deployment cannot be rolled back.
- During the rollback, connections will drop because the current configuration will be cleared.

Procedure

Step 1 At the Firewall Threat Defense CLI, roll back to the previous configuration.

configure policy rollback

After the rollback, the Firewall Threat Defense notifies the Firewall Management Center that the rollback was completed successfully. In the Firewall Management Center, the deployment screen will show a banner stating that the configuration was rolled back.

Note
If the rollback failed and the Firewall Management Center management is restored, refer to https://www.cisco.com/c/en/us/support/docs/security/firepower-ngfw-virtual/215258-troubleshooting-firepower-threat-defense.html for common deployment problems. In some cases, the rollback can fail after the Firewall Management Center management access is restored; in this case, you can resolve the Firewall Management Center configuration issues, and redeploy from the Firewall Management Center.

Example:

For the Firewall Threat Defense that uses a data interface for manager access:

```
> configure policy rollback

The last deployment to this FTD was on June 1, 2020 and its status was Successful.
Do you want to continue [Y/N]?

Y

Rolling back complete configuration on the FTD. This will take time.
.....................
Policy rollback was successful on the FTD.
Configuration has been reverted back to transaction id:
Following is the rollback summary:
.....................
.....................
>
```

Step 2 Check that the management connection was reestablished.

In Firewall Management Center, check the management connection status on the Connection Status page. Navigate to **Devices** > **Device Management** and then navigate to **Management** area under the **Devices** tab. Then in the **Manager Access - Configuration Details** screen, click **Connection Status**.

At the Firewall Threat Defense CLI, enter the **sftunnel-status-brief** command to view the management connection status.

If it takes more than 10 minutes to reestablish the connection, you should troubleshoot the connection. See Troubleshoot Management Connectivity on a Data Interface, on page 190.

Troubleshoot Management Connectivity on a Data Interface

When you use a data interface for manager access instead of using the dedicated Management interface, you must be careful about changing the interface and network settings for the Firewall Threat Defense in the Firewall Management Center so you do not disrupt the connection. If you change the management interface type after you add the Firewall Threat Defense to the Firewall Management Center (from data to Management, or from Management to data), if the interfaces and network settings are not configured correctly, you can lose management connectivity.

This topic helps you troubleshoot the loss of management connectivity.

View management connection status

In the Firewall Management Center, check the management connection status on the **Devices** > **Device Management** page.

At the Firewall Threat Defense CLI, enter the **sftunnel-status-brief** command to view the management connection status. You can also use **sftunnel-status** to view more complete information.

See the following sample output for a connection that is down; there is no peer channel "connected to" information, nor heartbeat information shown:

```
> sftunnel-status-brief
PEER:10.10.17.202
Registration: Completed.
Connection to peer '10.10.17.202' Attempted at Mon Jun 15 09:21:57 2020 UTC
Last disconnect time : Mon Jun 15 09:19:09 2020 UTC
Last disconnect reason : Both control and event channel connections with peer went down
```

See the following sample output for a connection that is up, with peer channel and heartbeat information shown:

```
> sftunnel-status-brief
PEER:10.10.17.202
Peer channel Channel-A is valid type (CONTROL), using 'eth0', connected to '10.10.17.202'
 via '10.10.17.222'
Peer channel Channel-B is valid type (EVENT), using 'eth0', connected to '10.10.17.202'
 via '10.10.17.222'
Registration: Completed.
IPv4 Connection to peer '10.10.17.202' Start Time: Wed Jun 10 14:27:12 2020 UTC
Heartbeat Send Time: Mon Jun 15 09:02:08 2020 UTC
Heartbeat Received Time: Mon Jun 15 09:02:16 2020 UTC
```

View the Firewall Threat Defense network information

At the Firewall Threat Defense CLI, view the Management and manager access data interface network settings:

show network

```
> show network
===============[ System Information ]===============
Hostname                : FTD-4
Domains                 : cisco.com
DNS Servers             : 72.163.47.11
DNS from router         : enabled
Management port         : 8305
IPv4 Default route
  Gateway               : data-interfaces

==================[ management0 ]===================
Admin State             : enabled
Admin Speed             : 1gbps
Operation Speed         : 1gbps
Link                    : up
Channels                : Management & Events
Mode                    : Non-Autonegotiation
MDI/MDIX                : Auto/MDIX
MTU                     : 1500
MAC Address             : 68:87:C6:A6:54:80
```

```
----------------------[ IPv4 ]----------------------
Configuration             : Manual
Address                   : 10.89.5.4
Netmask                   : 255.255.255.192
Gateway                   : 169.254.1.1
----------------------[ IPv6 ]----------------------
Configuration             : Disabled

===============[ Proxy Information ]================
State                     : Disabled
Authentication            : Disabled

======[ System Information - Data Interfaces ]======
DNS Servers               : 72.163.47.11
Interfaces                : Ethernet1/1

==================[ Ethernet1/1 ]===================
State                     : Enabled
Link                      : Up
Name                      : outside
MTU                       : 1500
MAC Address               : 68:87:C6:A6:54:A4
----------------------[ IPv4 ]----------------------
Configuration             : Manual
Address                   : 10.89.5.6
Netmask                   : 255.255.255.192
Gateway                   : 10.89.5.1
----------------------[ IPv6 ]----------------------
Configuration             : Disabled
```

Check that the Firewall Threat Defense registered with the Firewall Management Center

At the Firewall Threat Defense CLI, check that the Firewall Management Center registration was completed. Note that this command will not show the *current* status of the management connection.

show managers

```
> show managers
Type                      : Manager
Host                      : 16a3893c-caa7-11ee-8436-0925c06e7608DONTRESOLVE
Display name              : manager-1707852946.80444
Version                   : 7.6.0 (Build 1385)
Identifier                : a904b8b2-ca9a-11ee-a583-5e804c16b2fd
Registration              : Completed
Management type           : Configuration and analytics
```

Ping the Firewall Management Center

At the Firewall Threat Defense CLI, use the following command to ping the Firewall Management Center from the data interfaces:

ping *fmc_ip*

At the Firewall Threat Defense CLI, use the following command to ping the Firewall Management Center from the Management interface, which should route over the backplane to the data interfaces:

ping system *fmc_ip*

Capture packets on the Firewall Threat Defense internal interface

At the Firewall Threat Defense CLI, capture packets on the internal backplane interface (nlp_int_tap) to see if management packets are being sent:

capture *name* **interface nlp_int_tap trace detail match ip any any**

show capture*name* **trace detail**

Check the internal interface status, statistics, and packet count

At the Firewall Threat Defense CLI, see information about the internal backplane interface, nlp_int_tap:

show interface detail

```
> show interface detail
[...]
Interface Internal-Data0/1 "nlp_int_tap", is up, line protocol is up
  Hardware is en_vtun rev00, BW Unknown Speed-Capability, DLY 1000 usec
 (Full-duplex), (1000 Mbps)
 Input flow control is unsupported, output flow control is unsupported
 MAC address 0000.0100.0001, MTU 1500
 IP address 169.254.1.1, subnet mask 255.255.255.248
 37 packets input, 2822 bytes, 0 no buffer
 Received 0 broadcasts, 0 runts, 0 giants
 0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
 0 pause input, 0 resume input
 0 L2 decode drops
 5 packets output, 370 bytes, 0 underruns
 0 pause output, 0 resume output
 0 output errors, 0 collisions, 0 interface resets
 0 late collisions, 0 deferred
 0 input reset drops, 0 output reset drops
 input queue (blocks free curr/low): hardware (0/0)
 output queue (blocks free curr/low): hardware (0/0)
  Traffic Statistics for "nlp_int_tap":
 37 packets input, 2304 bytes
 5 packets output, 300 bytes
 37 packets dropped
      1 minute input rate 0 pkts/sec,  0 bytes/sec
      1 minute output rate 0 pkts/sec,  0 bytes/sec
      1 minute drop rate, 0 pkts/sec
      5 minute input rate 0 pkts/sec,  0 bytes/sec
      5 minute output rate 0 pkts/sec,  0 bytes/sec
      5 minute drop rate, 0 pkts/sec
  Control Point Interface States:
 Interface number is 14
 Interface config status is active
 Interface state is active
```

Check routing and NAT

At the Firewall Threat Defense CLI, check that the default route (S*) was added and that internal NAT rules exist for the Management interface (nlp_int_tap).

show route

```
> show route

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, V - VPN
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, + - replicated route
       SI - Static InterVRF
```

```
Gateway of last resort is 10.89.5.1 to network 0.0.0.0

S*       0.0.0.0 0.0.0.0 [1/0] via 10.89.5.1, outside
C        10.89.5.0 255.255.255.192 is directly connected, outside
L        10.89.5.29 255.255.255.255 is directly connected, outside

>
```

show nat

```
> show nat

Auto NAT Policies (Section 2)
1 (nlp_int_tap) to (outside) source static nlp_server_0_sftunnel_intf3 interface  service
 tcp 8305 8305
    translate_hits = 0, untranslate_hits = 6
2 (nlp_int_tap) to (outside) source static nlp_server_0_ssh_intf3 interface  service
tcp ssh ssh
    translate_hits = 0, untranslate_hits = 73
3 (nlp_int_tap) to (outside) source static nlp_server_0_sftunnel_ipv6_intf3 interface
ipv6  service tcp 8305 8305
    translate_hits = 0, untranslate_hits = 0
4 (nlp_int_tap) to (outside) source dynamic nlp_client_0_intf3 interface
    translate_hits = 174, untranslate_hits = 0
5 (nlp_int_tap) to (outside) source dynamic nlp_client_0_ipv6_intf3 interface ipv6
    translate_hits = 0, untranslate_hits = 0
>
```

Check other settings

See the following commands to check that all other settings are present. You can also see many of these commands on the Firewall Management Center's **Devices** > **Device Management** page.

show running-config sftunnel

```
> show running-config sftunnel
sftunnel interface outside
sftunnel port 8305
```

show running-config ip-client

```
> show running-config ip-client
ip-client outside
```

show conn address *fmc_ip*

```
> show conn address 10.89.5.35
5 in use, 16 most used
Inspect Snort:
        preserve-connection: 0 enabled, 0 in effect, 0 most enabled, 0 most in effect

TCP nlp_int_tap  10.89.5.29(169.254.1.2):51231 outside  10.89.5.35:8305, idle 0:00:04,
 bytes 86684, flags UxIO
TCP nlp_int_tap  10.89.5.29(169.254.1.2):8305 outside  10.89.5.35:52019, idle 0:00:02,
 bytes 1630834, flags UIO
>
```

Check for a successful DDNS update

At the Firewall Threat Defense CLI, check for a successful DDNS update:

debug ddns

```
> debug ddns
DDNS update request = /v3/update?hostname=domain.example.org&myip=209.165.200.225
Successfully updated the DDNS sever with current IP addresses
DDNS: Another update completed, outstanding = 0
DDNS: IDB SB total = 0
```

If the update failed, use the **debug http** and **debug ssl** commands. For certificate validation failures, check that the root certificates are installed on the device:

show crypto ca certificates *trustpoint_name*

To check the DDNS operation:

show ddns update interface *fmc_access_ifc_name*

```
> show ddns update interface outside

Dynamic DNS Update on outside:
    Update Method Name Update Destination
    RBD_DDNS not available

Last Update attempted on 04:11:58.083 UTC Thu Jun 11 2020
Status : Success
FQDN : domain.example.org
IP addresses : 209.165.200.225
```

Check Firewall Management Center log files

See https://cisco.com/go/fmc-reg-error.

View Inventory Details

The **Inventory Details** section of the **Device** page shows chassis details such as the CPU and memory.

Figure 86: Inventory Details

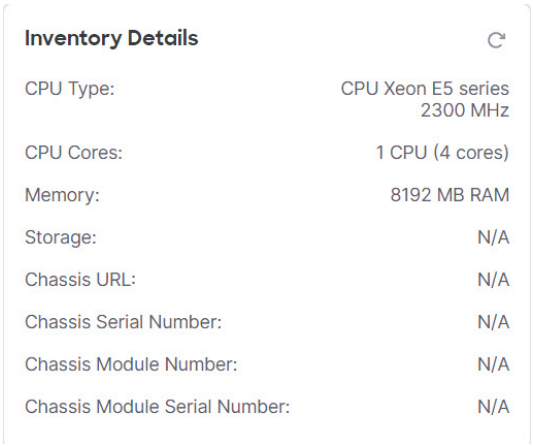

To update information, click **Refresh** ( ).

Edit Applied Policies

The **Applied Policies** section of the **Device** page displays the following policies applied to your firewall:

Figure 87: Applied Policies

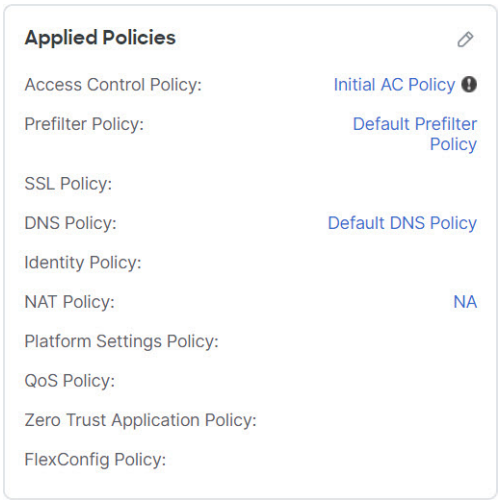

For policies with links, you can click the link to view the policy.

For the Access Control Policy, view the **Access Policy Information for Troubleshooting** dialog box by clicking the **Exclamation** ( ) icon. This dialog box shows how access rules are expanded into access control entries (ACEs).

Figure 88: Access Policy Information for Troubleshooting

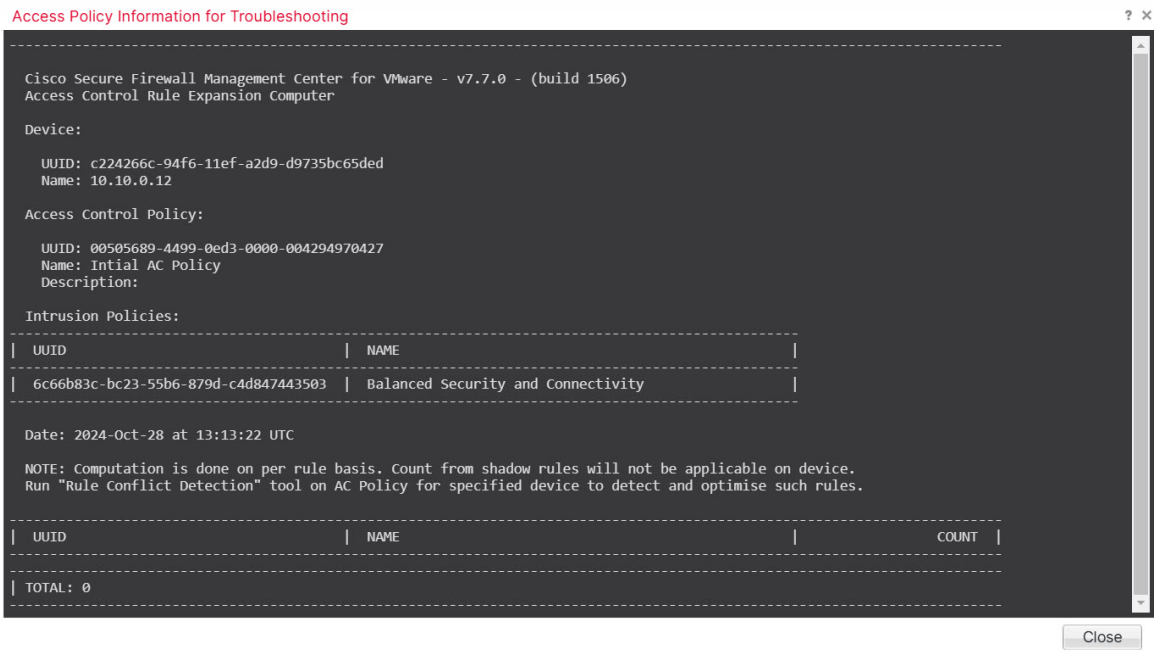

You can assign policies to an individual device from the **Device Management** page.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to assign policies, click **Edit** ( ).

Step 3 Click **Device**.

Step 4 In the **Applied Policies** section, click **Edit** ( ).

Figure 89: Policy Assignments

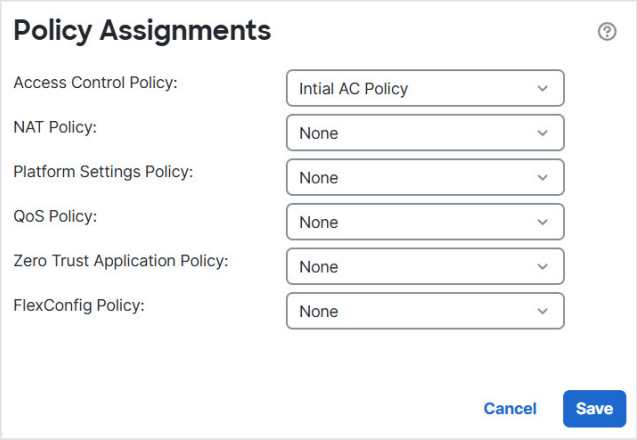

Step 5 For each policy type, choose a policy from the drop-down menu. Only existing policies are listed.

Step 6 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Edit Advanced Settings

The **Advanced Settings** section of the **Device** page displays a table of advanced configuration settings, as described below. You can edit any of these settings.

Table 8: Advanced Section Table Fields

Field	Description
Application Bypass	The state of Automatic Application Bypass on the device.
Bypass Threshold	The Automatic Application Bypass threshold, in milliseconds.

Field	Description
Object Group Search	The state of object group search on the device. While operating, the FTD device expands access control rules into multiple access control list entries based on the contents of any network or interface objects used in the access rule. You can reduce the memory required to search access control rules by enabling object group search. With object group search enabled, the system does not expand network or interface objects, but instead searches access rules for matches based on those group definitions. Object group search does not impact how your access rules are defined or how they appear in Firepower Management Center. It impacts only how the device interprets and processes them while matching connections to access control rules. **Note** By default, the **Object Group Search** is enabled when you add threat defense for the first time in the management center.
Interface Object Optimization	The state of interface object optimization on the device. During deployment, interface groups and security zones used in the access control and prefilter policies generate separate rules for each source/destination interface pair. If you enable interface object optimization, the system will instead deploy a single rule per access control/prefilter rule, which can simplify the device configuration and improve deployment performance. If you select this option, also select the **Object Group Search** option to reduce memory usage on the device.

The following topics explain how to edit the advanced device settings.

Note For information about the Transfer Packets setting, see Edit General Settings, on page 127.

Configure Automatic Application Bypass

Automatic Application Bypass (AAB) allows packets to bypass detection if Snort is down or, for a Classic device, if a packet takes too long to process. AAB causes Snort to restart within ten minutes of the failure, and generates troubleshooting data that can be analyzed to investigate the cause of the Snort failure.

Caution AAB activation partially restarts the Snort process, which temporarily interrupts the inspection of a few packets. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

See the following behavior:

Firewall Threat Defense Behavior: If Snort is down, then AAB is triggered after the specified timer duration. If Snort is up, then AAB is never triggered, even if packet processing exceeds the configured timer.

Classic Device Behavior: AAB limits the time allowed to process packets through an interface. You balance packet processing delays with your network's tolerance for packet latency.

The feature functions with any deployment; however, it is most valuable in inline deployments.

Typically, you use Rule Latency Thresholding in the intrusion policy to fast-path packets after the latency threshold value is exceeded. Rule Latency Thresholding does not shut down the engine or generate troubleshooting data.

If detection is bypassed, the device generates a health monitoring alert.

By default the AAB is disabled; to enable AAB follow the steps described.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to edit advanced device settings, click **Edit** ( ).

Step 3 Click **Device**, then click **Edit** ( ) in the **Advanced Settings** section.

Step 4 Check **Automatic Application Bypass**.

Step 5 Enter a **Bypass Threshold** from 250 ms to 60,000 ms. The default setting is 3000 milliseconds (ms).

Step 6 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Configure Object Group Search

While operating, the Firewall Threat Defense device expands access control rules into multiple access control list entries based on the contents of any network or interface objects used in the access rule. You can reduce the memory required to search access control rules by enabling object group search. With object group search enabled, the system does not expand network or interface objects, but instead searches access rules for matches based on those group definitions. Object group search does not impact how your access rules are defined or how they appear in Firewall Management Center. It impacts only how the device interprets and processes them while matching connections to access control rules.

Enabling object group search reduces memory requirements for access control policies that include network or interface objects. However, it is important to note that object group search might also decrease rule lookup performance and thus increase CPU utilization. You should balance the CPU impact against the reduced memory requirements for your specific access control policy. In most cases, enabling object group search provides a net operational improvement.

By default, the object group search is enabled for the threat defense devices that are added for the first time in the Firewall Management Center. In the case of upgraded devices, if the device is configured with disabled object group search, then you need to manually enable it. You can enable it on one device at a time; you cannot enable it globally. We recommend that you enable it on any device to which you deploy access rules that use network or interface objects.

Note If you enable object group search and then configure and operate the device for a while, be aware that subsequently disabling the feature might lead to undesirable results. When you disable object group search, your existing access control rules will be expanded in the device's running configuration. If the expansion requires more memory than is available on the device, your device can be left in an inconsistent state and you might see a performance impact. If your device is operating normally, you should not disable object group search once you have enabled it.

Before you begin

- Model Support—Threat Defense
- We recommend that you also enable transactional commit on each device. From the device CLI, enter the **asp rule-engine transactional-commit access-group** command.
- Changing this setting can be disruptive to system operation while the device recompiles the ACLs. We recommend that you change this setting during a maintenance window.
- You can use FlexConfig to configure the **object-group-search threshold** command to enable a threshold to help prevent performance degradation. When operating with a threshold, for each connection, both the source and destination IP addresses are matched against network objects. If the number of objects matched by the source address times the number matched by the destination address exceeds 10,000, the connection is dropped. Configure your rules to prevent an excessive number of matches.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the Firewall Threat Defense device where you want to configure the rule, click the **Edit** (✎).

Step 3 Click the **Device** tab, then click the **Edit** (✎) in the **Advanced Settings** section.

Step 4 Check **Object Group Search**.

Step 5 To have object group search work on interface objects in addition to network objects, check **Interface Object Optimization**.

If you do not select **Interface Object Optimization**, the system deploys separate rules for each source/interface pair, rather that use the security zones and interface groups used in the rules. This means the interface groups are not available for object group search processing.

Step 6 Click **Save**.

Configure Interface Object Optimization

During deployment, interface groups and security zones used in the access control and prefilter policies generate separate rules for each source/destination interface pair. If you enable interface object optimization, the system will instead deploy a single rule per access control/prefilter rule, which can simplify the device configuration and improve deployment performance. If you select this option, also select the **Object Group Search** option to reduce memory usage on the device.

Interface object optimization is disabled by default. You can enable it on one device at a time; you cannot enable it globally.

Note If you disable interface object optimization, your existing access control rules will be deployed without using interface objects, which might make deployment take longer. In addition, if object group search is enabled, its benefits will not apply to interface objects, and you might see expansion in the access control rules in the device's running configuration. If the expansion requires more memory than is available on the device, your device can be left in an inconsistent state and you might see a performance impact.

Before you begin

Model Support—Threat Defense

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the Firewall Threat Defense device where you want to configure the rule, click the **Edit** ( ).

Step 3 Click the **Device** tab, then click **Edit** ( ) in the **Advanced Settings** section.

Step 4 Check **Interface Object Optimization**.

Step 5 Click **Save**.

Edit Deployment Settings

The **Deployment Settings** section of the **Device** page displays the information described in the table below.

Figure 90: Deployment Settings

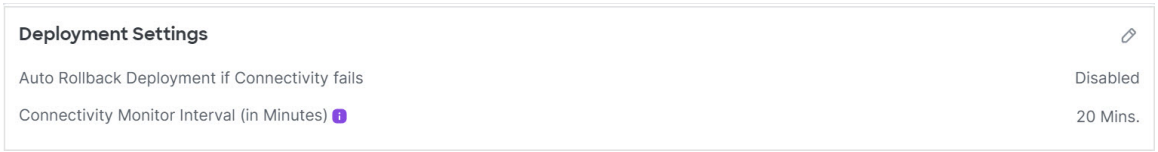

Table 9: Deployment Settings

Field	Description
Auto Rollback Deployment if Connectivity Fails	Enabled or Disabled. You can enable auto rollback if the management connection fails as a result of the deployment; specifically if you use data for management center access, and then you misconfigure the data interface.
Connectivity Monitor Interval (in Minutes)	Shows the amount of time to wait before rolling back the configuration.

You can set deployment settings from the **Device Management** page. Deployment settings include enabling auto rollback of the deployment if the management connection fails as a result of the deployment; specifically if you use data for management center access, and then you misconfigure the data interface. You can alternatively manually roll back the configuration using the **configure policy rollback** command (see Manually Roll Back the Configuration if the Firewall Management Center Loses Connectivity, on page 189).

See the following guidelines:

- Only the previous deployment is available locally on the Firewall Threat Defense; you cannot roll back to any earlier deployments.
- Rollback is supported for high availability but not supported for clustering deployments.
- The rollback only affects configurations that you can set in the Firewall Management Center. For example, the rollback does not affect any local configuration related to the dedicated Management interface, which you can only configure at the Firewall Threat Defense CLI. Note that if you changed data interface settings after the last Firewall Management Center deployment using the **configure network management-data-interface** command, and then you use the rollback command, those settings will not be preserved; they will roll back to the last-deployed Firewall Management Center settings.
- UCAPL/CC mode cannot be rolled back.
- Out-of-band SCEP certificate data that was updated during the previous deployment cannot be rolled back.
- During the rollback, connections will drop because the current configuration will be cleared.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device where you want to assign policies, click **Edit** ( ).

Step 3 Click **Device**.

Step 4 In the **Deployment Settings** section, click **Edit** ( ).

Figure 91: Deployment Settings

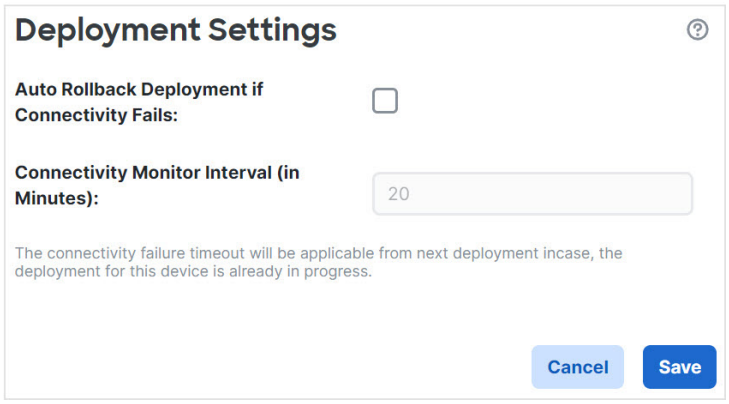

Step 5 Check **Auto Rollback Deployment if Connectivity Fails** to enable auto rollback.

Step 6 Set the **Connectivity Monitor Interval (in Minutes)** to set the amount of time to wait before rolling back the configuration. The default is 20 minutes.

Step 7 If a rollback occurs, see the following for next steps.

- If the auto rollback was successful, you see a success message instructing you to do a full deployment.
- You can also go to the **Deploy**and then **Advanced Deploy** screen and click the **Preview** ( ) icon to view the parts of the configuration that were rolled back (see Deploy Configuration Changes, on page 259). Click **Show Rollback Changes** to view the changes, and **Hide Rollback Changes** to hide the changes.

Figure 92: Rollback Changes

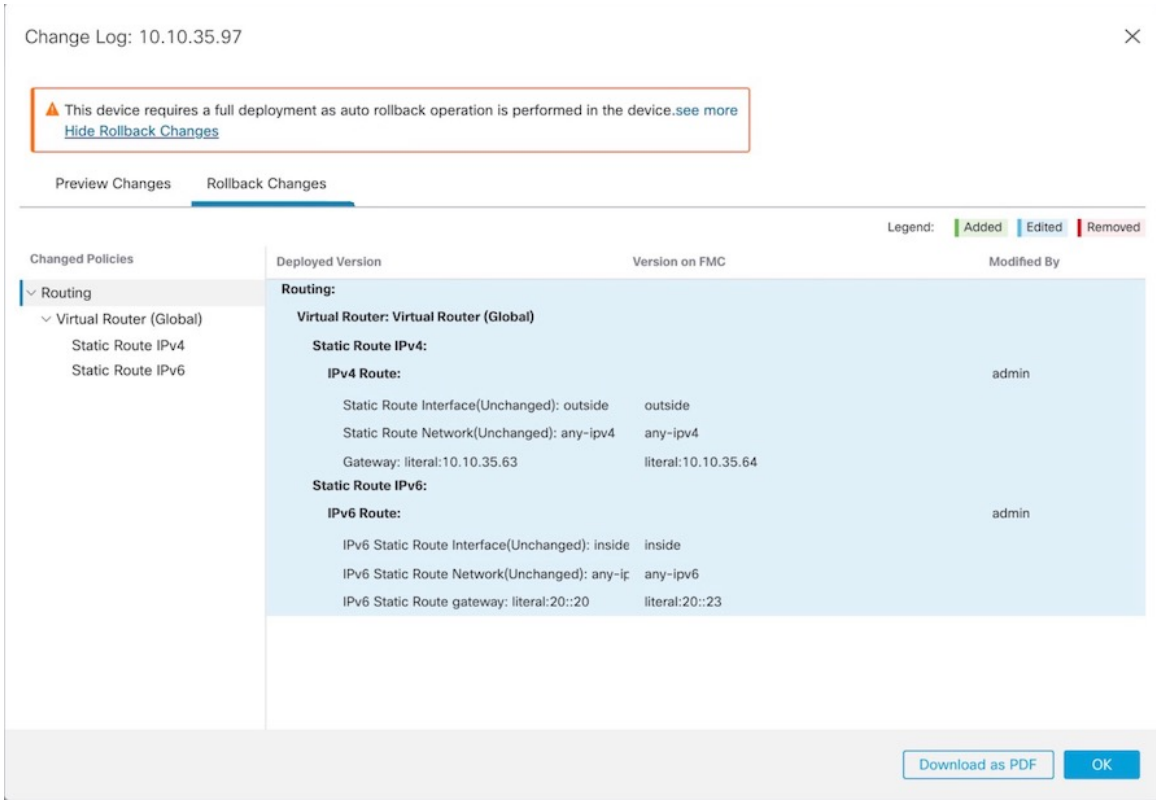

- In the Deployment History Preview, you can view the rollback changes. See View Deployment History, on page 267.

Step 8 Check that the management connection was reestablished.

In Firewall Management Center, check the management connection status on the connection status page. Navigate to **Devices** > **Device Management** and then in **Management** area under the **Devices** tab, click **Connection Status** to view the **Connection Status** page.

At the Firewall Threat Defense CLI, enter the **sftunnel-status-brief** command to view the management connection status.

If it takes more than 10 minutes to reestablish the connection, you should troubleshoot the connection. See Troubleshoot Management Connectivity on a Data Interface, on page 190.

Edit Cluster Health Monitor Settings

The **Cluster Health Monitor Settings** section of the **Cluster** page displays the settings described in the table below.

Figure 93: Cluster Health Monitor Settings

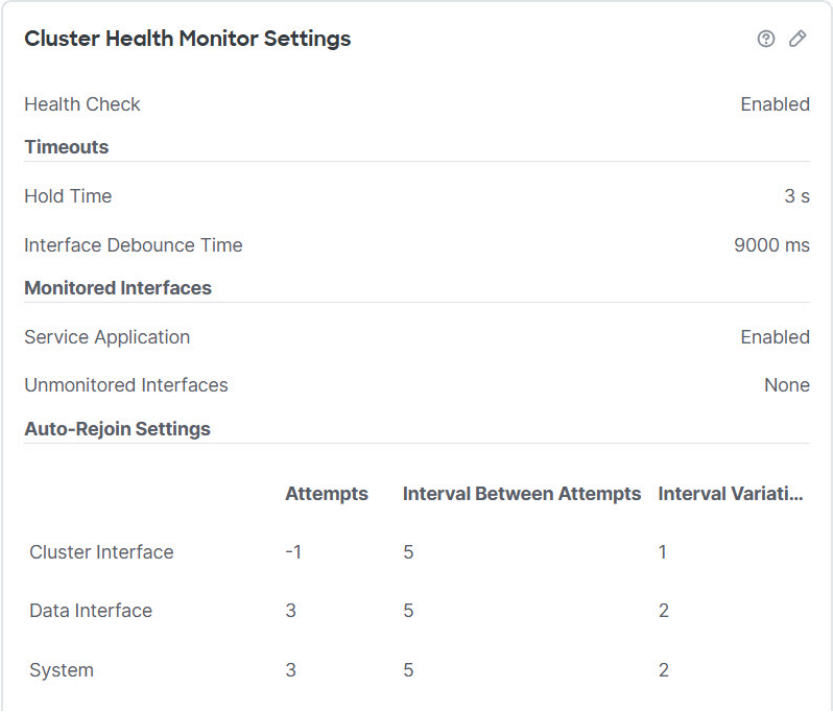

Table 10: Cluster Health Monitor Settings Section Table Fields

Field	Description
Timeouts	
Hold Time	Between .3 and 45 seconds; The default is 3 seconds. To determine node system health, the cluster nodes send heartbeat messages on the cluster control link to other nodes. If a node does not receive any heartbeat messages from a peer node within the hold time period, the peer node is considered unresponsive or dead.
Interface Debounce Time	Between 300 and 9000 ms. The default is 500 ms. The interface debounce time is the amount of time before the node considers an interface to be failed, and the node is removed from the cluster.

Field	Description
Monitored Interfaces	The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The amount of time before the node removes a member from the cluster depends on the type of interface and whether the node is an established node or is joining the cluster.
Service Application	Shows whether the Snort and disk-full processes are monitored.
Unmonitored Interfaces	Shows unmonitored interfaces.
Auto-Rejoin Settings	
Cluster Interface	Shows the auto-rejoin settings after a cluster control link failure.
Attempts	Between -1 and 65535. The default is -1 (unlimited). Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 1x the interval duration. Defines if the interval duration increases at each attempt.
Data Interfaces	Shows the auto-rejoin settings after a data interface failure.
Attempts	Between -1 and 65535. The default is 3. Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 2x the interval duration. Defines if the interval duration increases at each attempt.
System	Shows the auto-rejoin settings after internal errors. Internal failures include: application sync timeout; inconsistent application statuses; and so on.
Attempts	Between -1 and 65535. The default is 3. Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 2x the interval duration. Defines if the interval duration increases at each attempt.

Note If you disable the system health check, fields that do not apply when the system health check is disabled will not show.

You can change these settings from this section.

You can monitor any port-channel ID, single physical interface ID, as well as the Snort and disk-full processes. Health monitoring is not performed on VLAN subinterfaces or virtual interfaces such as VNIs or BVIs. You cannot configure monitoring for the cluster control link; it is always monitored.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the cluster you want to modify, click **Edit** (✎).

Step 3 Click **Cluster**.

Step 4 In the **Cluster Health Monitor Settings** section, click **Edit** (✎).

Step 5 Disable the system health check by clicking the **Health Check** slider .

Figure 94: Disable the System Health Check

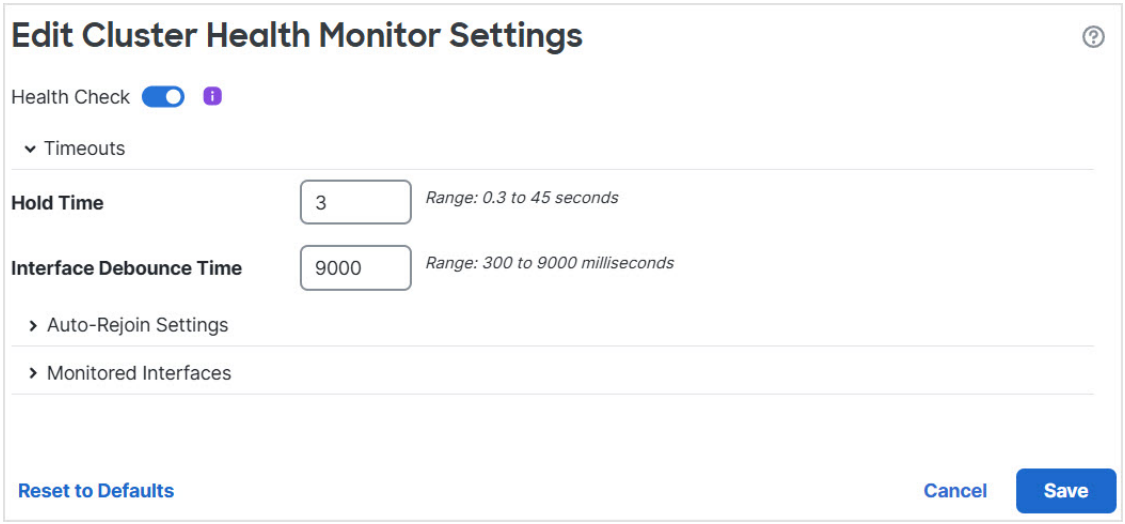

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 6 Configure the hold time and interface debounce time.

- **Hold Time**—Set the hold time to determine the amount of time between node heartbeat status messages, between .3 and 45 seconds; The default is 3 seconds.
- **Interface Debounce Time**—Set the debounce time between 300 and 9000 ms. The default is 500 ms. Lower values allow for faster detection of interface failures. Note that configuring a lower debounce time increases the chances of false-positives. When an interface status update occurs, the node waits the number of milliseconds specified before marking the interface as failed, and the node is removed from the cluster. In the case of an EtherChannel that transitions from a down state to an up state (for example, the switch reloaded, or the switch enabled an EtherChannel), a longer debounce time can prevent the interface from appearing to be failed on a cluster node just because another cluster node was faster at bundling the ports.

Step 7 Customize the auto-rejoin cluster settings after a health check failure.

Figure 95: Configure Auto-Rejoin Settings

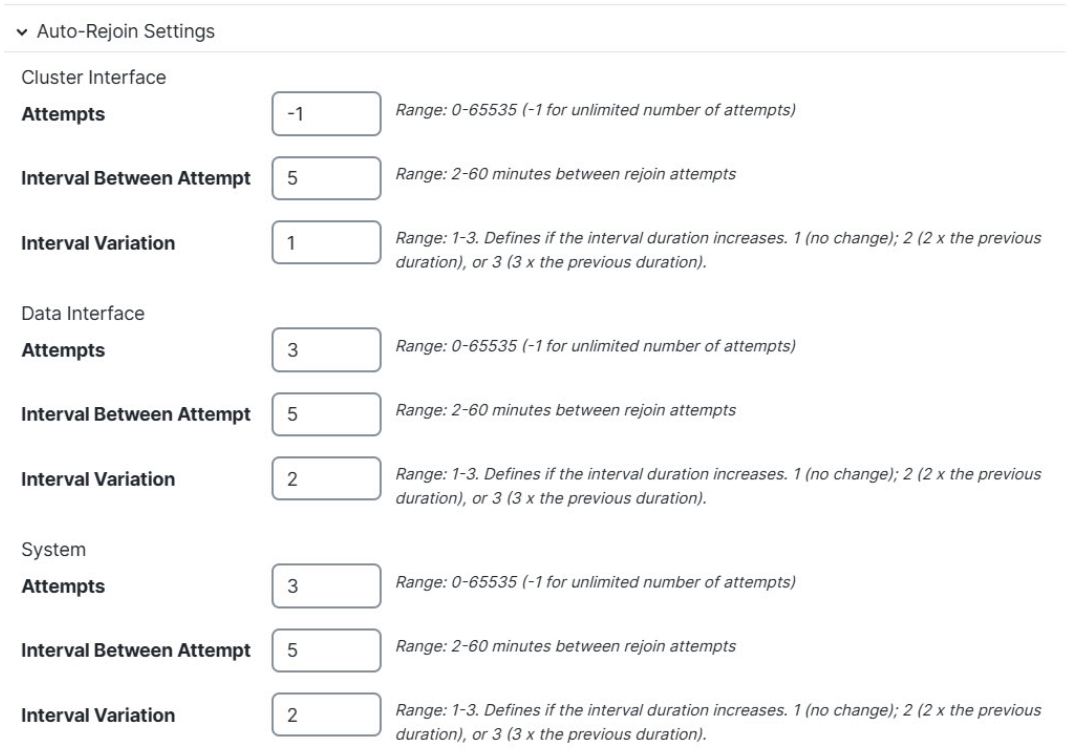

Set the following values for the **Cluster Interface**, **Data Interface**, and **System** (internal failures include: application sync timeout; inconsistent application statuses; and so on):

- **Attempts**—Sets the number of rejoin attempts, between -1 and 65535. **0** disables auto-rejoining. The default for the **Cluster Interface** is -1 (unlimited). The default for the **Data Interface** and **System** is 3.
- **Interval Between Attempts**—Defines the interval duration in minutes between rejoin attempts, between 2 and 60. The default value is 5 minutes. The maximum total time that the node attempts to rejoin the cluster is limited to 14400 minutes (10 days) from the time of last failure.
- **Interval Variation**—Defines if the interval duration increases. Set the value between 1 and 3: **1** (no change); **2** (2 x the previous duration), or **3** (3 x the previous duration). For example, if you set the interval duration to 5 minutes, and set the variation to 2, then the first attempt is after 5 minutes; the 2nd attempt is 10 minutes (2 x 5); the 3rd attempt 20 minutes (2 x 10), and so on. The default value is **1** for the **Cluster Interface** and **2** for the **Data Interface** and **System**.

Step 8 Configure monitored interfaces by moving interfaces in the **Monitored Interfaces** or **Unmonitored Interfaces** window. You can also check or uncheck **Enable Service Application Monitoring** to enable or disable monitoring of the Snort and disk-full processes.

Figure 96: Configure Monitored Interfaces

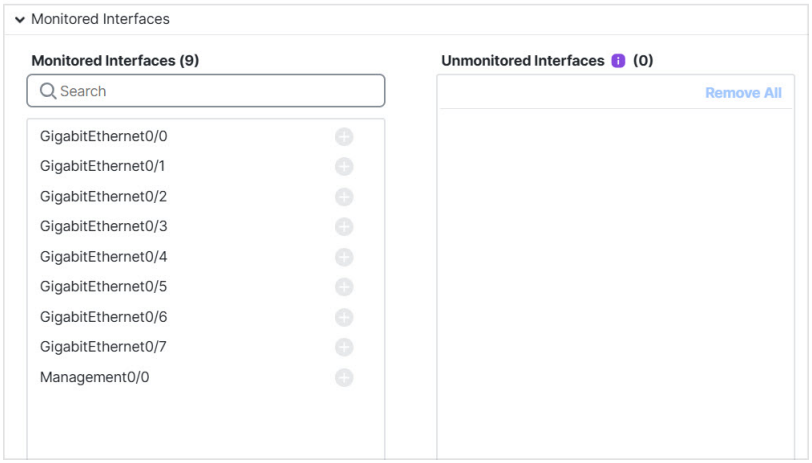

The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The amount of time before the node removes a member from the cluster depends on the type of interface and whether the node is an established node or is joining the cluster. Health check is enabled by default for all interfaces and for the Snort and disk-full processes.

You might want to disable health monitoring of non-essential interfaces.

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 9 Click **Save**.

Step 10 Deploy configuration changes; see Deploy Configuration Changes, on page 259.

History for Device Settings

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Recovery-config mode for emergency on-device configuration and out-of-band configuration detection on the Firewall Management Center	7.7.0	7.7.0	If you lose the management connection to your device, you can make select configuration changes directly at the device CLI to: • Restore the management connection if you are using a data interface for manager access • Make select policy changes that can't wait until the connection is restored After the management connection is restored, the Firewall Management Center will detect the configuration changes on the device. It does not automatically update the device configuration in the Firewall Management Center; you must view the configuration differences, acknowledge that the device configuration is different, and then manually make the same changes in the Firewall Management Center before you deploy. New/modified diagnostic CLI (**system support diagnostic-cli**) command: **configure recovery-config** New/modified screens: **Devices** > **Device Management** > **Device** > **Health** > **Out of Band Status**
High availability is supported with redundant manager access data interfaces	7.7.0	7.7.0	You can now use redundant manager access data interfaces with high availability.
View CLI output for a device or device cluster.	7.4.1	Any	You can view a set of pre-defined CLI outputs that can help you troubleshoot the device or cluster. You can also enter any **show** command and see the output. New/modified screens: **Devices** > **Device Management** > **Cluster** > **General**
Troubleshooting file generation and download available from Device and Cluster pages.	7.4.1	7.4.1	You can generate and download troubleshooting files for each device on the Device page and also for all cluster nodes on the Cluster page. For a cluster, you can download all files as a single compressed file. You can also include cluster logs for the cluster for cluster nodes. You can alternatively trigger file generation from the **Devices** > **Device Management** > **More** > **Troubleshoot Files** menu. New/modified screens: • **Devices** > **Device Management** > **Device** > **General** • **Devices** > **Device Management** > **Cluster** > **General**

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Cluster health monitor settings.	7.3.0	Any	You can now edit cluster health monitor settings. New/modified screens: **Devices** > **Device Management** > **Cluster** > **Cluster Health Monitor Settings** **Note** If you previously configured these settings using FlexConfig, be sure to remove the FlexConfig configuration before you deploy. Otherwise the FlexConfig configuration will overwrite the management center configuration.
Redundant manager access data interface.	7.3.0	7.3.0	When you use a data interface for manager access, you can configure a secondary data interface to take over management functions if the primary interface goes down. The device uses SLA monitoring to track the viability of the static routes and an ECMP zone that contains both interfaces so management traffic can use both interfaces. New/modified screens: • **Devices** > **Device Management** > **Device** > **Management** • **Devices** > **Device Management** > **Device** > **Interfaces** > **Manager Access**
Policy rollback support for high availability devices.	7.2.0	7.2.0	The **configure policy rollback** command is supported for high availability devices.
Auto rollback of a deployment that causes a loss of management connectivity.	7.2.0	7.2.0	You can now enable auto rollback of the configuration if a deployment causes the management connection between the management center and the threat defense to go down. Previously, you could only manually rollback a configuration using the **configure policy rollback** command. New/modified screens: • **Devices** > **Device Management** > **Device** > **Deployment Settings** • **Deploy** > **Advanced Deploy** > **Preview** • **Deploy** > **Deployment History** > **Preview**
Object group search is enabled by default for access control rules.	7.2.0	7.2.0	The **Object Group Search** setting is enabled by default for managed devices starting with Version 7.2.0. This option is in the **Advanced Settings** section when editing device settings on the Device Management page.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Import and export device configurations.	7.1.0	7.1.0	You can export the device-specific configuration, and you can then import the saved configuration for the same device in the following use cases: • Moving the device to a different FMC. • Restore an old configuration. • Reregistering a device. New/modified screens: **Devices** > **Device Management** > **Device** > **General**
Update the FMC IP address on FTD.	6.7.0	6.7.0	If you change the FMC IP address, you can now use the FTD CLI to update the device. New/modified commands: **configure manager edit**

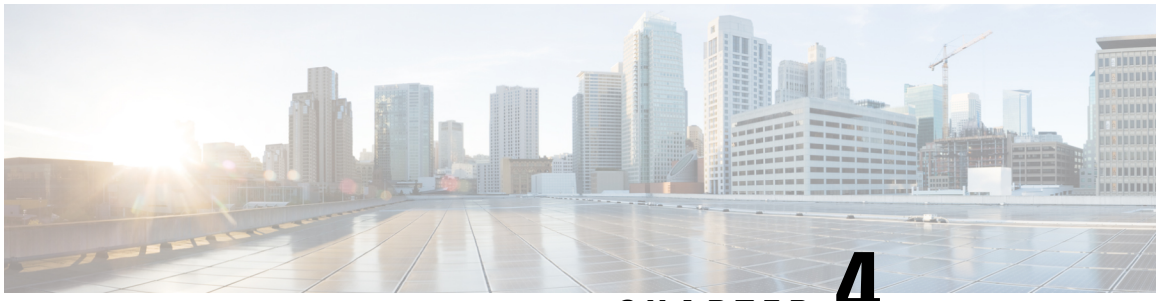

CHAPTER 4

Users

Managed devices include a default **admin** account for CLI access. This chapter discusses how to create custom user accounts.

- About Users, on page 213
- Requirements and Prerequisites for User Accounts for Devices, on page 214
- Guidelines and Limitations for User Accounts for Devices, on page 215
- Add an Internal User at the CLI, on page 215
- Configure External Authentication for the Firewall Threat Defense, on page 217
- Troubleshooting LDAP Authentication Connections, on page 229
- History for Users, on page 231

About Users

You can add custom user accounts on managed devices, either as internal users or as external users on a LDAP or RADIUS server. Each managed device maintains separate user accounts. For example, when you add a user to the Firewall Management Center, that user only has access to the Firewall Management Center; you cannot then use that username to log directly into a managed device. You must separately add a user on the managed device.

Internal and External Users

Managed devices support two types of users:

- Internal user—The device checks a local database for user authentication.
- External user—If the user is not present in the local database, the system queries an external LDAP or RADIUS authentication server.

CLI Access

Firepower devices include a Firepower CLI that runs on top of Linux. You can create internal users on devices using the CLI. You can establish external users on Firewall Threat Defense devices using the Firewall Management Center.

Caution Users with CLI Config level access can access the Linux shell using the **expert** command, and obtain `sudoers` privileges in the Linux shell, which can present a security risk. For system security reasons, we strongly recommend:

- Only use the Linux shell under TAC supervision or when explicitly instructed by Firepower user documentation.
- Make sure that you restrict the list of users with CLI access appropriately.
- When granting CLI access privileges, restrict the list of users with Config level access.
- Do not add users directly in the Linux shell; only use the procedures in this chapter.
- Do not access Firepower devices using CLI expert mode unless directed by Cisco TAC or by explicit instructions in the Firepower user documentation.

CLI User Roles

On managed devices, user access to commands in the CLI depends on the role you assign.

None

The user cannot log into the device on the command line.

Config

The user can access all commands, including configuration commands. Exercise caution in assigning this level of access to users.

Basic

The user can access non-configuration commands only. Allowed commands are dig, ping, and traceroute. Only internal users and Firewall Threat Defense external RADIUS users support the Basic role.

Requirements and Prerequisites for User Accounts for Devices

Model Support

- Threat Defense—Internal and external users

Supported Domains

Any

User Roles

Configure external users—Admin FMC user

Configure internal users—Config CLI user

Guidelines and Limitations for User Accounts for Devices

Usernames

- You cannot add the same username for both internal and external users. If the external server uses a duplicate username, the deployment to the device fails.
- The username must be Linux-valid:
  - Maximum 32 alphanumeric characters, plus hyphen (-) and underscore (_)
  - All lowercase
  - Cannot start with hyphen (-); cannot be all numbers; cannot include a period (.), at sign (@), or slash (/)

Defaults

All devices include an **admin** user as a local user account; you cannot delete the **admin** user. The default initial password is **Admin123**; the system forces you to change this during the initialization process. The agent stores the modified credentials in the Windows registry and uses the same encryption as the Firewall Management Center, which is Cipher Block Chaining (CBC). See the getting started guide for your model for more information about system initialization.

Number of User Accounts

You can create a maximum of 43 user accounts for the Firepower 1000.

Add an Internal User at the CLI

Use the CLI to create internal users on the Firewall Threat Defense.

Procedure

Step 1 Log into the device CLI using an account with Config privileges.

The **admin** user account has the required privileges, but any account with Config privileges will work. You can use an SSH session or the Console port.

For certain Firewall Threat Defense models, the Console port puts you into the FXOS CLI. Use the **connect ftd** command to get to the Firewall Threat Defense CLI.

Step 2 Create the user account.

configure user add *username* {**basic** | **config**}

- *username*—Sets the username. The username must be Linux-valid:
  - Maximum 32 alphanumeric characters, plus hyphen (-) and underscore (_)

- All lowercase
- Cannot start with hyphen (-); cannot be all numbers; cannot include a period (.), at sign (@), or slash (/)

- **basic**—Gives the user basic access. This role does not allow the user to enter configuration commands. Allowed commands are dig, ping, and traceroute.
- **config**—Gives the user configuration access. This role gives the user full administrator rights to all commands.

Example:

The following example adds a user account named johncrichton with Config access rights. The password is not shown as you type it.

```
> configure user add johncrichton config
Enter new password for user johncrichton: newpassword
Confirm new password for user johncrichton: newpassword
> show user
Login              UID   Auth Access  Enabled Reset    Exp Warn  Str Lock Max
admin             1000  Local Config  Enabled    No  Never  N/A  Dis   No N/A
johncrichton      1001  Local Config  Enabled    No  Never  N/A  Dis   No   5
```

Note
Tell users they can change their own passwords using the **configure password** command.

Step 3 (Optional) Adjust the characteristics of the account to meet your security requirements.

You can use the following commands to change the default account behavior.

- **configure user aging** *username max_days warn_days*

  Sets an expiration date for the user's password. Specify the maximum number of days for the password to be valid followed by the number of days before expiration the user will be warned about the upcoming expiration. Both values are 1 to 9999, but the warning days must be less than the maximum days. When you create the account, there is no expiration date for the password.

- **configure user forcereset** *username*

  Forces the user to change the password on the next login.

- **configure user maxfailedlogins** *username number*

  Sets the maximum number of consecutive failed logins you will allow before locking the account, from 1 to 9999. Use the **configure user unlock** command to unlock accounts. The default for new accounts is 5 consecutive failed logins.

- **configure user minpasswdlen** *username number*

  Sets a minimum password length, which can be from 1 to 127.

- **configure user strengthcheck** *username* {**enable** | **disable**}

  Enables or disables password strength checking, which requires a user to meet specific password criteria when changing their password. When a user's password expires or if the **configure user forcereset** command is used, this requirement is automatically enabled the next time the user logs in.

Step 4 Manage user accounts as necessary.

Users can get locked out of their accounts, or you might need to remove accounts or fix other issues. Use the following commands to manage the user accounts on the system.

- **configure user access** *username* {**basic** | **config**}

  Changes the privileges for a user account.

- **configure user delete** *username*

  Deletes the specified account.

- **configure user disable** *username*

  Disables the specified account without deleting it. The user cannot log in until you enable the account.

- **configure user enable** *username*

  Enables the specified account.

- **configure user password** *username*

  Changes the password for the specified user. Users should normally change their own password using the **configure password** command.

  Caution
  Do not use the Linux **passwd** command in Expert Mode to change the admin user password. This command can cause file system corruption. Only use the regular Firewall Threat Defense CLI **configure user password admin** command (if you are not admin) or **configure password** command (if you are admin). If you don't know the password and can't log in at all, see the password recovery procedure.

- **configure user unlock** *username*

  Unlocks a user account that was locked due to exceeding the maximum number of consecutive failed login attempts.

Configure External Authentication for the Firewall Threat Defense

To enable external authentication for Firewall Threat Defense devices, you need to add one or more external authentication objects.

About External Authentication for the Firewall Threat Defense

When you enable external authentication for Firewall Threat Defense users, the Firewall Threat Defense verifies the user credentials with an LDAP or RADIUS server as specified in an *external authentication object*.

External authentication objects can be used by the Firewall Management Center and Firewall Threat Defense devices. You can share the same object between the different appliance/device types or create separate objects. For the Firewall Threat Defense, you can only activate one external authentication object in the platform settings that you deploy to the devices.

Note The timeout range is different for the Firewall Threat Defense and the Firewall Management Center, so if you share an object, be sure not to exceed the Firewall Threat Defense's smaller timeout range (1-30 seconds for LDAP, and 1-300 seconds for RADIUS). If you set the timeout to a higher value, the Firewall Threat Defense external authentication configuration will not work.

Only a subset of fields in the external authentication object are used for Firewall Threat Defense SSH access. If you fill in additional fields, they are ignored. If you also use this object for other device types, those fields will be used.

LDAP users always have Config privileges. RADIUS users can be defined as either Config or Basic users.

You can either define users on the RADIUS server (with the Service-Type attribute), or you can pre-define the user list in the external authentication object. For LDAP, you can specify a filter to match CLI users on the LDAP server.

Note Users with CLI access can gain Linux shell access with the **expert** command. Linux shell users can obtain root privileges, which can present a security risk. Make sure that you:

- Restrict the list of users with Linux shell access.
- Do not create Linux shell users.

About LDAP

The Lightweight Directory Access Protocol (LDAP) allows you to set up a directory on your network that organizes objects, such as user credentials, in a centralized location. Multiple applications can then access those credentials and the information used to describe them. If you ever need to change a user's credentials, you can change them in one place.

Microsoft has announced that Active Directory servers will start enforcing LDAP binding and LDAP signing in 2020. Microsoft is making these a requirement because when using default settings, an elevation of privilege vulnerability exists in Microsoft Windows that could allow a man-in-the-middle attacker to successfully forward an authentication request to a Windows LDAP server. For more information, see 2020 LDAP channel binding and LDAP signing requirement for Windows on the Microsoft support site.

If you have not done so already, we recommend you start using TLS/SSL encryption to authenticate with an Active Directory server.

About RADIUS

Remote Authentication Dial In User Service (RADIUS) is an authentication protocol used to authenticate, authorize, and account for user access to network resources. You can create an authentication object for any RADIUS server that conforms to RFC 2865.

Secure Firewall devices support the use of SecurID tokens. When you configure authentication by a server using SecurID, users authenticated against that server append the SecurID token to the end of their SecurID PIN and use that as their password when they log in. You do not need to configure anything extra on the Secure Firewall device to support SecurID.

Add an LDAP External Authentication Object for Firewall Threat Defense

Add an LDAP server to support external users for Firewall Threat Defense management.

In a multidomain deployment, external authentication objects are only available in the domain in which they are created.

Sharing External Authentication Objects

External LDAP objects can be used by the Firewall Management Center and Firewall Threat Defense devices. You can share the same object between the Firewall Management Center and devices or create separate objects. Be sure that both the Firewall Threat Defense and the Firewall Management Center can reach the LDAP server, even if you are not sharing the object. The Firewall Management Center is essential to retrieving the user list and downloading it to the device.

Note For LDAP, the timeout range is different for the Firewall Threat Defense and the Firewall Management Center, so if you share an object, be sure not to exceed the Firewall Threat Defense's smaller timeout range (1-30 seconds). If you set the timeout to a higher value, the deployment to the Firewall Threat Defense will fail.

Firewall Threat Defense Supported Fields

Only a subset of fields in the LDAP object are used for Firewall Threat Defense SSH access. If you fill in additional fields, they are ignored. If you also use this object for the Firewall Management Center, those fields will be used. This procedure only covers the supported fields for the Firewall Threat Defense. For other fields, see Add an LDAP External Authentication Object for the Firewall Management Center.

Usernames

Usernames must be Linux-valid usernames and be lower-case only, using alphanumeric characters plus period (.) or hyphen (-). Other special characters such as at sign (@) and slash (/) are not supported. You cannot add the **admin** user for external authentication. You can only add external users (as part of the External Authentication object) in the Firewall Management Center; you cannot add them at the CLI. Note that internal users can only be added at the CLI, not in the Firewall Management Center.

If you previously configured the same username for an internal user using the **configure user add** command, the Firewall Threat Defense first checks the password against the internal user, and if that fails, it checks the LDAP server. Note that you cannot later add an internal user with the same name as an external user; only pre-existing internal users are supported.

Privilege Level

LDAP users always have Config privileges.

Before you begin

You must specify DNS server(s) for domain name lookup on your device. Even if you specify an IP address and not a hostname for the LDAP server on this procedure, the LDAP server may return a URI for authentication that can include a hostname. A DNS lookup is required to resolve the hostname. See Modify Firewall Threat Defense Management Interfaces at the CLI, on page 158 to add DNS servers.

Procedure

Step 1 Choose **System** (⚙) > **Users**.

Step 2 Click the **External Authentication** tab.

Step 3 Click (+) **Add External Authentication Object**.

Step 4 Set the **Authentication Method** to **LDAP**.

Step 5 Enter a **Name** and optional **Description**.

Step 6 Choose a **Server Type** from the drop-down list.

Step 7 For the **Primary Server**, enter a **Host Name/IP Address**.

If you are using a certificate to connect via TLS or SSL, the host name in the certificate must match the host name used in this field. In addition, IPv6 addresses are not supported for encrypted connections.

Step 8 (Optional) Change the **Port** from the default.

Step 9 (Optional) Enter the **Backup Server** parameters.

Step 10 Enter **LDAP-Specific Parameters**.

a) Enter the **Base DN** for the LDAP directory you want to access. For example, to authenticate names in the Security organization at the Example company, enter `ou=security,dc=example,dc=com`. Alternatively click **Fetch DNs**, and choose the appropriate base distinguished name from the drop-down list.

b) (Optional) Enter the **Base Filter**. For example, if the user objects in a directory tree have a `physicalDeliveryOfficeName` attribute and users in the New York office have an attribute value of `NewYork` for that attribute, to retrieve only users in the New York office, enter `(physicalDeliveryOfficeName=NewYork)`.

c) Enter a **User Name** for a user who has sufficient credentials to browse the LDAP server. For example, if you are connecting to an OpenLDAP server where user objects have a `uid` attribute, and the object for the administrator in the Security division at your example company has a `uid` value of `NetworkAdmin`, you might enter `uid=NetworkAdmin,ou=security,dc=example,dc=com`.

d) Enter the user password in the **Password** and the **Confirm Password** fields.

e) (Optional) Click **Show Advanced Options** to configure the following advanced options.

- **Encryption**—Click **None**, **TLS**, or **SSL**.

  If you change the encryption method after specifying a port, you reset the port to the default value for that method. For **None** or **TLS**, the port resets to the default value of 389. If you choose SSL encryption, the port resets to 636.

- **SSL Certificate Upload Path**—For SSL or TLS encryption, you must choose a certificate by clicking **Choose File**.

  If you previously uploaded a certificate and want to replace it, upload the new certificate and redeploy the configuration to your devices to copy over the new certificate.

  Note
  TLS encryption requires a certificate on all platforms. For SSL, the Firewall Threat Defense also requires a certificate. For other platforms, SSL does not require a certificate. However, we recommend that you *always* upload a certificate for SSL to prevent man-in-the-middle attacks.

- (Not Used) **User Name Template**—Not used by the Firewall Threat Defense.

- **Timeout (Seconds)**—Enter the number of seconds before rolling over to the backup connection, between 1 and 30. The default is 30.

  Note
  The timeout range is different for the Firewall Threat Defense and the Firewall Management Center, so if you share an object, be sure not to exceed the Firewall Threat Defense's smaller timeout range (1-30 seconds). If you set the timeout to a higher value, the Firewall Threat Defense LDAP configuration will not work.

Step 11 Configure **Attribute Mapping** to retrieve users based on an attribute.

- Enter a **UI Access Attribute**. **Note**: This field is not used for device CLI access; however, it is a required field, so you need to enter a value. You can just enter the same value that you enter for the **CLI Access Attribute**.
- Set the **CLI Access Attribute** if you want to use a CLI access attribute other than the user distinguished type. For example, on a Microsoft Active Directory Server, use the `sAMAccountName` CLI access attribute to retrieve CLI access users by typing `sAMAccountName`.

  Note
  Users with CLI access can gain Linux shell access with the **expert** command. Linux shell users can obtain root privileges, which can present a security risk. Make sure that you restrict the list of users with CLI or Linux shell access.

  Note
  Deploying an external authentication object that allows a large number of users with CLI access may cause deployments to time out and fail while waiting for the users to be created.

Step 12 Set the **CLI Access Filter**.

Choose one of the following methods:

- To use the same filter you specified when configuring authentication settings, check the check box of **Same as Base Filter**.
- To retrieve administrative user entries based on attribute value, enter the attribute name, a comparison operator, and the attribute value you want to use as a filter, enclosed in parentheses. For example, if all network administrators have a `manager` attribute which has an attribute value of `shell`, you can set a base filter of `(manager=shell)`.

The usernames must be Linux-valid:

- Maximum 32 alphanumeric characters, plus hyphen (-) and underscore (_)
- All lowercase
- Cannot start with hyphen (-); cannot be all numbers; cannot include a period (.), at sign (@), or slash (/)

Note
If you previously configured the same username for an internal user, the Firewall Threat Defense first checks the password against the internal user, and if that fails, it checks the LDAP server. Note that you cannot later add an internal user with the same name as an external user; only pre-existing internal users are supported.

Step 13 Click **Save**.

Step 14 Enable use of this server. See External Authentication, on page 952.

Step 15 If you later add or delete users on the LDAP server, you must redeploy the Platform Settings on managed devices. The Firewall Management Center redownloads the user list and deploys it to the device. See Deploy Configuration Changes, on page 259.

Examples

Basic Example

The following figures illustrate a basic configuration of an LDAP login authentication object for a Microsoft Active Directory Server. The LDAP server in this example has an IP address of 10.11.3.4. The connection uses port 389 for access.

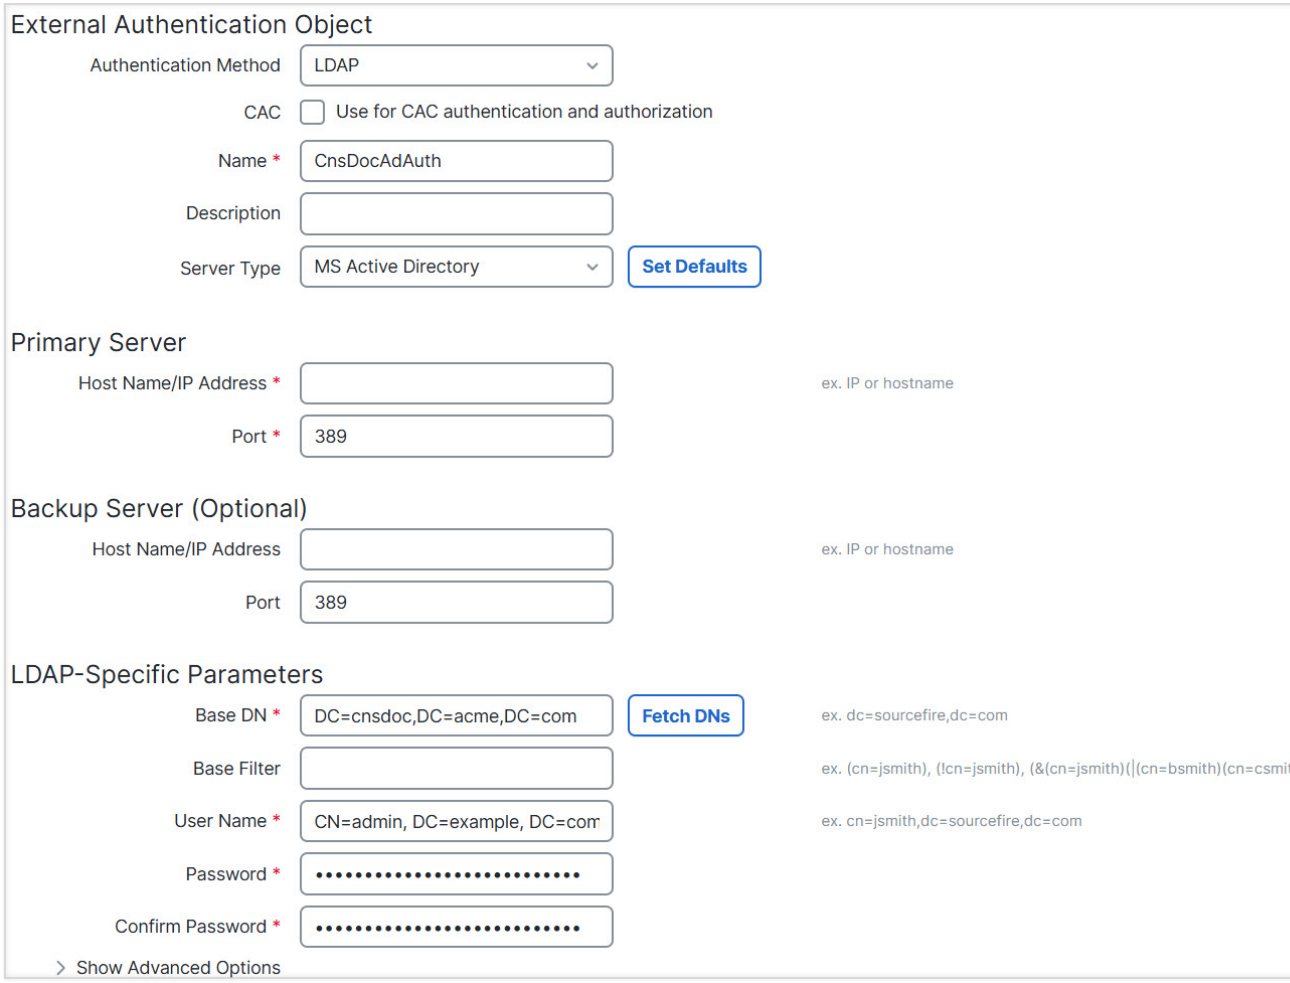

This example shows a connection using a base distinguished name of `OU=security,DC=it,DC=example,DC=com` for the security organization in the information technology domain of the Example company.

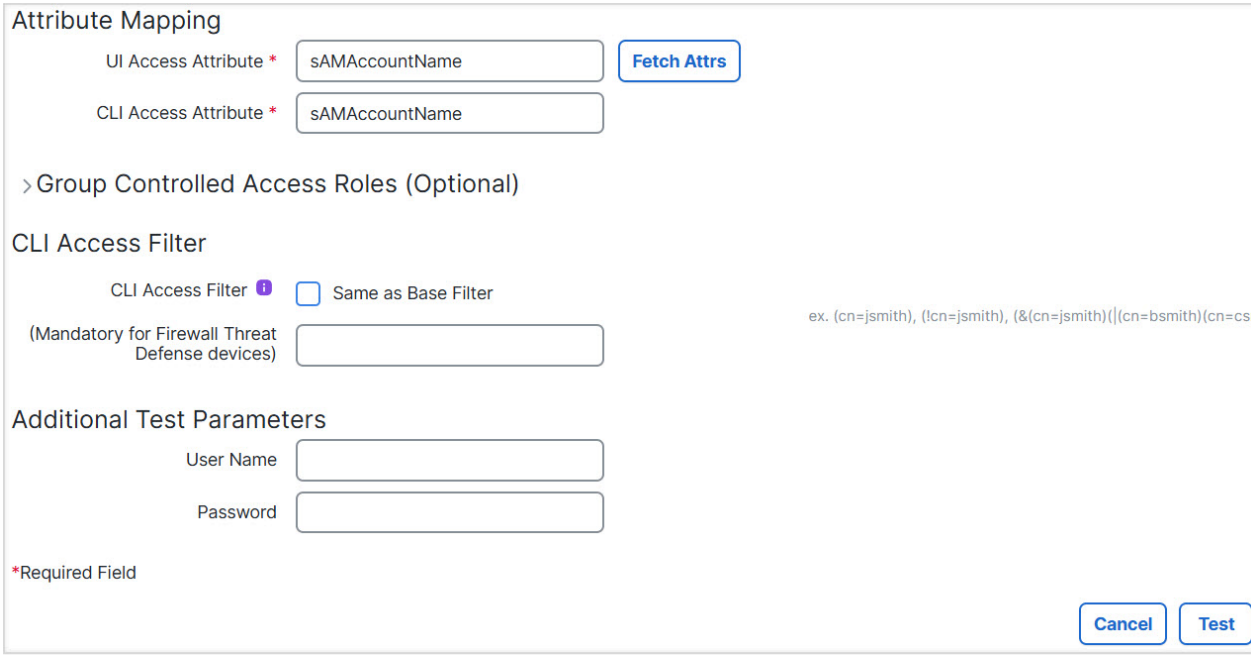

A **CLI Access Attribute** of `sAMAccountName` causes each `sAMAccountName` attribute to be checked for all objects in the directory for matches when a user logs into the Firewall Threat Defense.

Note that because no base filter is applied to this server, the Firewall Threat Defense checks attributes for all objects in the directory indicated by the base distinguished name. Connections to the server time out after the default time period (or the timeout period set on the LDAP server).

Advanced Example

This example illustrates an advanced configuration of an LDAP login authentication object for a Microsoft Active Directory Server. The LDAP server in this example has an IP address of 10.11.3.4. The connection uses port 636 for access.

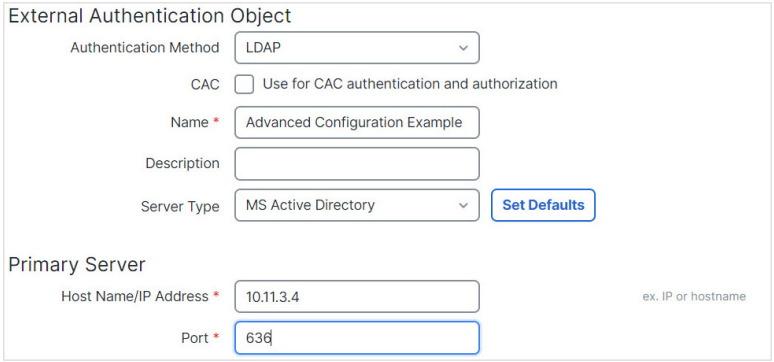

This example shows a connection using a base distinguished name of `OU=security,DC=it,DC=example,DC=com` for the security organization in the information technology domain of the Example company. However, note that this server has a base filter of `(cn=*smith)`. The filter restricts the users retrieved from the server to those with a common name ending in `smith`.

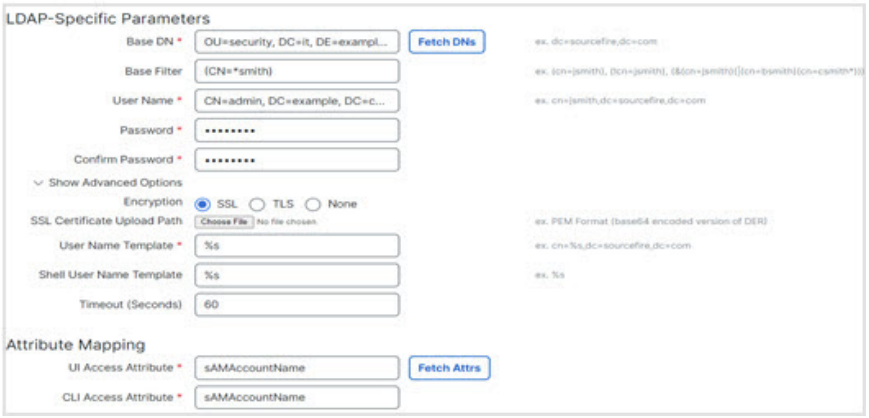

The connection to the server is encrypted using SSL and a certificate named `certificate.pem` is used for the connection. In addition, connections to the server time out after 60 seconds because of the **Timeout (Seconds)** setting.

Because this server is a Microsoft Active Directory server, it uses the `sAMAccountName` attribute to store user names rather than the `uid` attribute.

The **CLI Access Attribute** of `sAMAccountName` causes each `sAMAccountName` attribute to be checked for all objects in the directory for matches when a user logs into the Firewall Threat Defense.

In the following example, the CLI access filter is set to be the same as the base filter.

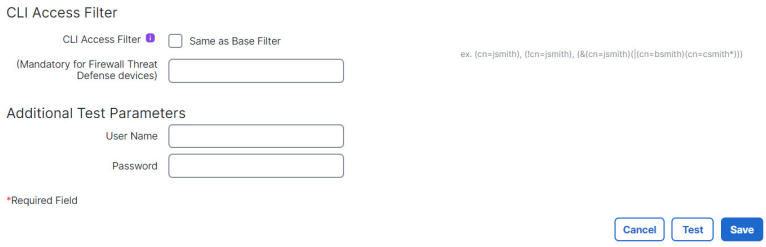

Add a RADIUS External Authentication Object for Firewall Threat Defense

Add a RADIUS server to support external users for the Firewall Threat Defense.

Sharing External Authentication Objects

You can share the same object between the Firewall Management Center and devices or create separate objects. Note that the Firewall Threat Defense supports defining users on the RADIUS server, while the Firewall Management Center requires you to predefine the user list in the external authentication object. You can choose to use the predefined list method for the Firewall Threat Defense, but if you want to define users on the RADIUS server, you must create separate objects for the Firewall Threat Defense and the Firewall Management Center.

Note The timeout range is different for the Firewall Threat Defense and the Firewall Management Center, so if you share an object, be sure not to exceed the Firewall Threat Defense's smaller timeout range (1-300 seconds). If you set the timeout to a higher value, the Firewall Threat Defense RADIUS configuration will not work.

Firewall Threat Defense Supported Fields

Only a subset of fields in the RADIUS object are used for Firewall Threat Defense SSH access. If you fill in additional fields, they are ignored. If you also use this object for the Firewall Management Center, those fields will be used. This procedure only covers the supported fields for the Firewall Threat Defense. For other fields, see Add a RADIUS External Authentication Object for Firewall Management Center in the Cisco Secure Firewall Management Center Administration Guide.

Usernames

You cannot add the **admin** user for external authentication. You can only add external users (as part of the External Authentication object) in the Firewall Management Center; you cannot add them at the CLI. Note that internal users can only be added at the CLI, not in the Firewall Management Center.

If you previously configured the same username for an internal user using the **configure user add** command, the Firewall Threat Defense first checks the password against the internal user, and if that fails, it checks the RADIUS server. Note that you cannot later add an internal user with the same name as an external user; only pre-existing internal users are supported. For users defined on the RADIUS server, be sure to set the privilege level to be the same as any internal users; otherwise you cannot log in using the external user password.

Procedure

Step 1 Define users on the RADIUS server using the Service-Type attribute.

The following are supported values for the Service-Type attribute:

- Administrator (6)—Provides Config access authorization to the CLI. These users can use all commands in the CLI.
- NAS Prompt (7) or any level other than 6—Provides Basic access authorization to the CLI. These users can use read-only commands, such as **show** commands, for monitoring and troubleshooting purposes.

The names must be Linux-valid usernames:

- Maximum 32 alphanumeric characters, plus hyphen (-) and underscore (_)
- All lowercase
- Cannot start with hyphen (-); cannot be all numbers; cannot include a sign (@) or slash (/)

Alternatively, you can predefine users in the external authentication object (see, Step 12, on page 226). To use the same RADIUS server for the Firewall Threat Defense and Firewall Management Center while using the Service-Type attribute method for the Firewall Threat Defense, create two external authentication objects that identify the same RADIUS server: one object includes the predefined **CLI Access Filter** users (for use with the Firewall Management Center), and the other object leaves the **CLI Access Filter** empty (for use with Firewall Threat Defenses).

Step 2 In the Firewall Management Center, choose **System** (⚙) > **Users**.

Step 3 Click **External Authentication**.

Step 4 Click (+)**Add External Authentication Object**.

Step 5 Set the **Authentication Method** to **RADIUS**.

Step 6 Enter a **Name** and optional **Description**.

Step 7 For the **Primary Server**, enter a **Host Name/IP Address**.

Only IPv4 is supported.

Note
If you are using a certificate to connect via TLS or SSL, the host name in the certificate must match the host name used in this field.

Step 8 (Optional) Change the **Port** from the default.

Step 9 Enter the **RADIUS Secret Key**.

Step 10 (Optional) Enter the **Backup Server** parameters.

Step 11 (Optional) Enter **RADIUS-Specific Parameters**.

a) Enter the **Timeout (Seconds)** in seconds before retrying the primary server, between 1 and 300. The default is 30.

Note
The timeout range is different for the Firewall Threat Defense and the Firewall Management Center, so if you share an object, be sure not to exceed the Firewall Threat Defense's smaller timeout range (1-300 seconds). If you set the timeout to a higher value, the Firewall Threat Defense RADIUS configuration will not work.

b) Enter the **Retries** before rolling over to the backup server. The default is 3.

Step 12 (Optional) Instead of using RADIUS-defined users (see, Step 1, on page 225), in the **CLI Access Filter** area **Administrator CLI Access User List** field, enter the user names that should have CLI access, separated by commas. For example, enter `jchrichton, aerynsun, rygel`.

You may want to use the **CLI Access Filter** method for Firewall Threat Defense so you can use the same external authentication object with Firewall Threat Defense and other platform types.

Note
If you want to use RADIUS-defined users, you must leave the **CLI Access Filter** empty.

Make sure that these usernames match usernames on the RADIUS server. The names must be Linux-valid usernames:

- Maximum 32 alphanumeric characters, plus hyphen (-) and underscore (_)
- All lowercase
- Cannot start with hyphen (-); cannot be all numbers; cannot include a period (.), at sign (@), or slash (/)

Note
Users with CLI access can gain Linux shell access with the **expert** command. Linux shell users can obtain root privileges, which can present a security risk. Make sure that you restrict the list of users with CLI or Linux shell access.

Note
Deploying an external authentication object that allows a large number of users with CLI access may cause deployments to time out and fail while waiting for the users to be created.

Step 13 (Optional) Click **Test** to test Firewall Management Center connectivity to the RADIUS server.

This function can only test Firewall Management Center connectivity to the RADIUS server; there is no test function for managed device connectivity to the RADIUS server.

Step 14 (Optional) You can also enter **Additional Test Parameters** to test user credentials for a user who should be able to authenticate: enter a **User Name** and **Password**, and then click **Test**.

Tip
If you mistype the name or password of the test user, the test fails even if the server configuration is correct. To verify that the server configuration is correct, click **Test** without entering user information in the **Additional Test Parameters** field first. If that succeeds, supply a user name and password to test with the specific user.

Example:

To test if you can retrieve the `JSmith` user credentials at the Example company, enter `JSmith` and the correct password.

Step 15 Click **Save**.

Step 16 Enable use of this server. See External Authentication, on page 952

Examples

Simple User Role Assignments

The following figure illustrates a sample RADIUS login authentication object for a server running Cisco Identity Services Engine (ISE) with an IP address of 10.10.10.98 on port 1812. No backup server is defined.

The following example shows RADIUS-specific parameters, including the timeout (30 seconds) and number of failed retries before the system attempts to contact the backup server, if any.

This example illustrates important aspects of RADIUS user role configuration:

Users `ewharton` and `gsand` are granted web interface Administrative access.

The user `cbronte` is granted web interface Maintenance User access.

The user `jausten` is granted web interface Security Analyst access.

The user `ewharton` can log into the device using a CLI account.

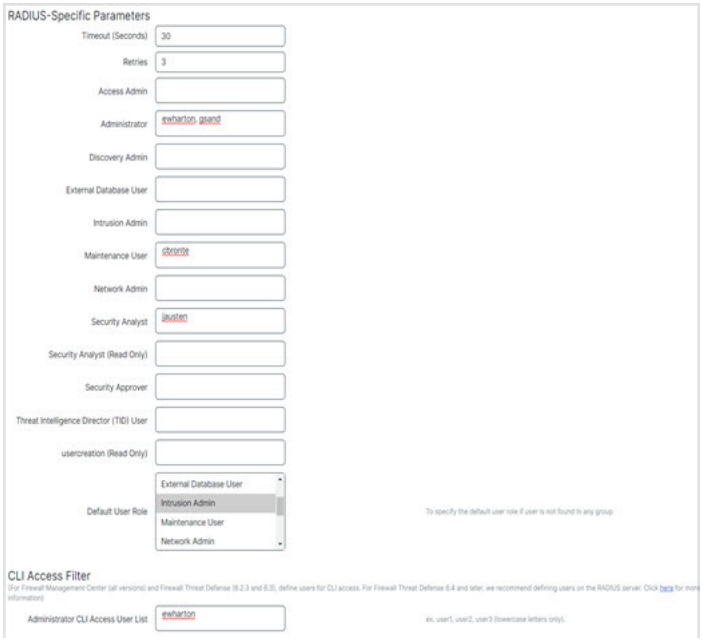

The following graphic depicts the role configuration for the example:

Roles for Users Matching an Attribute-Value Pair

You can use an attribute-value pair to identify users who should receive a particular user role. If the attribute you use is a custom attribute, you must define the custom attribute.

The following figure illustrates the role configuration and custom attribute definition in a sample RADIUS login authentication object for the same ISE server as in the previous example.

In this example, however, the `MS-RAS-Version` custom attribute is returned for one or more of the users because a Microsoft remote access server is in use. Note the `MS-RAS-Version` custom attribute is a string. In this example, all users logging in to RADIUS through a Microsoft v. 5.00 remote access server should receive the Security Analyst (Read Only) role, so you enter the attribute-value pair of `MS-RAS-Version=MSRASV5.00` in the **Security Analyst (Read Only)** field.

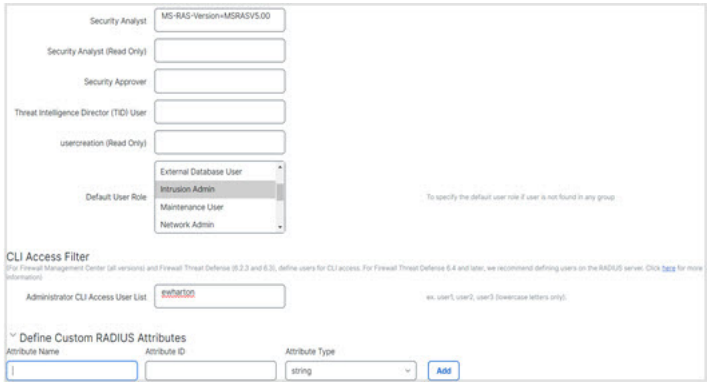

Enable External Authentication for Users on Firewall Threat Defense Devices

Enable External Authentication in the Threat Defense Platform Settings, and then deploy the settings to the managed devices. See External Authentication, on page 952 for more information.

Troubleshooting LDAP Authentication Connections

If you create an LDAP authentication object and it either does not succeed in connecting to the server you select or does not retrieve the list of users you want, you can tune the settings in the object.

If the connection fails when you test it, try the following suggestions to troubleshoot your configuration:

- Use the messages displayed at the top of the web interface screen and in the test output to determine which areas of the object are causing the issue.
- Check that the user name and password you used for the object are valid:
  - Check that you have the rights to browse to the directory indicated in your base-distinguished name by connecting to the LDAP server using a third-party LDAP browser.
  - Check that the user name is unique to the directory information tree for the LDAP server.
  - If you see an LDAP bind error 49 in the test output, the user binding for the user failed. Try authenticating to the server through a third-party application to see if the binding fails through that connection as well.
- Check that you have correctly identified the server:
  - Check that the server IP address or host name is correct.
  - Check that you have TCP/IP access from your local appliance to the authentication server where you want to connect.
  - Check that access to the server is not blocked by a firewall and that the port you have configured in the object is open.
  - If you are using a certificate to connect via TLS or SSL, the host name in the certificate must match the host name used for the server.
  - Check that you have not used an IPv6 address for the server connection if you are authenticating CLI access.
  - If you used server type defaults, check that you have the correct server type and click **Set Defaults** again to reset the default values.
- If you typed in your base-distinguished name, click **Fetch DNs** to retrieve all the available base distinguished names on the server, and select the name from the list.
- If you are using any filters, access attributes, or advanced settings, check that each is valid and typed correctly.
- If you are using any filters, access attributes, or advanced settings, try removing each setting and testing the object without it.

- If you are using a base filter or a CLI access filter, make sure that the filter is enclosed in parentheses and that you are using a valid comparison operator (maximum 450 characters, including the enclosing parentheses).
- To test a more restricted base filter, try setting it to the base distinguished name for the user to retrieve just that user.
- If you are using an encrypted connection:
  - Check that the name of the LDAP server in the certificate matches the host name that you use to connect.
  - Check that you have not used an IPv6 address with an encrypted server connection.
- If you are using a test user, make sure that the user name and password are typed correctly.
- If you are using a test user, remove the user credentials and test the object.
- Test the query that you are using by connecting to the LDAP server and using this syntax:

```
ldapsearch -x -b 'base_distinguished_name'
-h LDAPserver_ip_address -p port -v -D
'user_distinguished_name' -W 'base_filter'
```

  For example, if you are trying to connect to the security domain on `myrtle.example.com` using the `domainadmin@myrtle.example.com` user and a base filter of (`cn=*`), you could test the connection using this statement:

```
ldapsearch -x -b 'CN=security,DC=myrtle,DC=example,DC=com'
-h myrtle.example.com -p 389 -v -D
'domainadmin@myrtle.example.com' -W '(cn=*)'
```

If you can test your connection successfully but authentication does not work after you deploy a platform settings policy, check that authentication and the object you want to use are both enabled in the platform settings policy that is applied to the device.

If you connect successfully but want to adjust the list of users retrieved by your connection, you can add or change a base filter or CLI access filter or use a more restrictive or less restrictive base DN.

While authenticating a connection to Active Directory (AD) server, rarely the connection event log indicates blocked LDAP traffic although the connection to AD server is successful. This incorrect connection log occurs when the AD server sends a duplicate reset packet. The Firewall Threat Defense device identifies the second reset packet as part of a new connection request and logs the connection with Block action.

History for Users

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Limited user privileges for Threat Defense CLI Basic user.	7.7.0	7.7.0	The scope of the Threat Defense CLI Basic user privilege is now limited to the following commands: dig, ping, traceroute. If you have created users with the Basic privilege, evaluate whether you need to change them to the Config privilege. You can change a user's privilege level using the **configure user access** command. See: Cisco Secure Firewall Threat Defense Command Reference
Support for the Service-Type attribute for Firewall Threat Defense users defined on the RADIUS server.	6.4.0	Any	For RADIUS authentication of Firewall Threat Defense CLI users, you used to have to pre-define the usernames in the RADIUS external authentication object and manually make sure that the list matched usernames defined on the RADIUS server. You can now define CLI users on the RADIUS server using the Service-Type attribute and also define both Basic and Config user roles. To use this method, be sure to leave the shell access filter blank in the external authentication object. New/modified screens: **System** > **Users** > **External Authentication** > **Add External Authentication Object** > **Shell Access Filter**
External authentication for Firewall Threat Defense SSH access.	6.2.3	Any	You can now configure external authentication for SSH access to the Firewall Threat Defense using LDAP or RADIUS. New/modified screens: **Devices** > **Platform Settings** > **External Authentication**

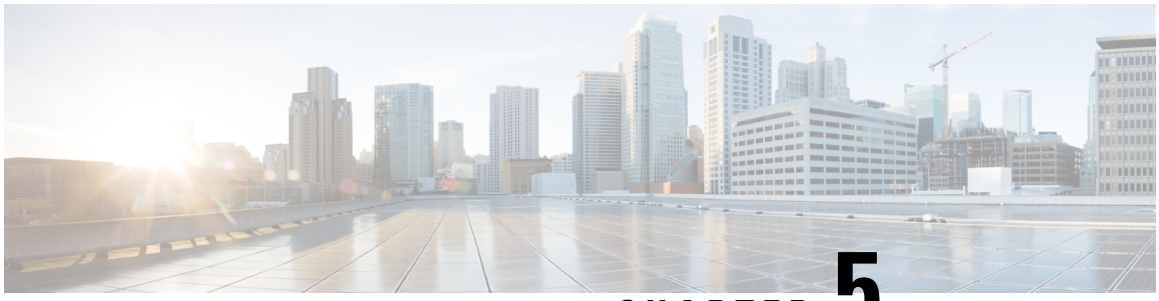

CHAPTER 5

Change Management

You can enable Change Management if your organization needs to implement more formal processes for configuration changes, including audit tracking and official approval before changes are deployed.

- About Change Management, on page 233
- Requirements and Prerequisites for Change Management, on page 237
- Guidelines and Limitations for Change Management, on page 237
- Enabling or Disabling Change Management, on page 238
- Managing Tickets, on page 239
- History for Change Management, on page 244

About Change Management

Some organizations need to implement a formal approach to deploying configuration changes. This might include more auditing, and an official approval process that must happen before configuration changes can be made to a device.

If your organization uses a more formal configuration change process, you can enable Change Management to enforce the process. With Change Management, administrators must open tickets before they can make configuration changes. Then, once the change is complete, they must submit the ticket for approval before they can deploy the proposed changes. This allows you to enforce your official approval process and ensure that the right employees make the final decisions.

When using Change Management, administrators can see their own changes within a ticket, but they cannot see changes anyone else has made within a ticket. Because a policy is locked once a user makes a change within a ticket, users should not be able to make interfering changes. However, users will not be able to make changes while another user has made a change that is pending approval.

Administrators can create multiple tickets so that a single ticket contains only logically-related policy changes. Tickets with a more limited scope are also easier to evaluate and approve quickly.

The following topics explain the Change Management Workflow and which policies and objects are subject to the ticketing and approval process.

How to Configure Devices in the Change Management Workflow

When you enable Change Management, users who configure devices need to change their approach slightly. Configuration specialists need to take the following approach when making configuration changes to supported policies and objects.

Procedure

Step 1 Create a ticket.

Step 2 Open the ticket.

Step 3 Make the configuration changes.

Note that the procedures explained in the online help and user guides assume that Change Management is not active, and omit any steps for creating, opening, or submitting tickets.

Step 4 Optionally, preview and validate the ticket to ensure the changes are complete and correct.

Step 5 Submit the ticket. At this point, the approver can either approve or reject the ticket.

- If the ticket is approved, deploy the changes.
- If the ticket is rejected, address the issues, and resubmit the ticket.

Creating Separate Approver and Configuration Roles

Some system-defined roles have permissions to modify (create/open/discard) and review (approve/reject) tickets:

- To both modify and review tickets:
  - Admin
  - Network Admin
- To modify tickets only:
  - Access Admin
  - Intrusion Admin
- To review tickets only:
  - Security Approver

If you need more granular roles to separate these activities due to your organizational requirements, you can create separate roles to ensure that ticket approval is assigned only to those users who have the organizational authority to approve changes. To create new user roles, go to **System** > **Users**, and select the **User Roles** tab..

Following are the permissions, in the **System** > **Change Management** folder, relevant to ticket usage and approval. Note that these permissions are available only after you enable Change Management.

- Modify Tickets—To create tickets (for yourself), to use tickets for configuration changes, and to discard tickets.
- Review Tickets—To approve or reject tickets.
- Both Modify and Review Tickets—To create tickets for yourself and others, use tickets, and approve/reject tickets. You can also take over tickets assigned to other users.

The approach you take depends on your precise requirements. For example:

- If your approvers should also be allowed to make configuration changes, you can simply assign them the system-defined roles, such as Administrator. Then, create custom configuration-only roles that include the same permissions but not the Review Tickets permission.
- If you need complete separation between approvers and those who make configuration changes, create custom roles for both, limiting the roles to either the Modify Tickets or the Review Tickets permission plus all other needed permissions for viewing or changing the supported policies and objects.

Policies and Objects that Support Change Management

If a policy or object supports the change management workflow, then creating, editing, or deleting the policy or object, including assigning a policy to a device, must be done in an open ticket.

Any action, policy, or object that does not support the change management workflow can be created, edited, or deleted, and so forth, without an open ticket. Even if a ticket is open, the changes made to unsupported policies are not included in the ticketed changes and are available for deployment immediately.

The following lists include the policies and objects that are supported. Anything not listed is unsupported.

Supported Policies

- Access Control, including rules, references to other policies, and inheritance settings. Cloning access control policies is not included within a ticket. The clone is immediately accepted and available to all users.
- Device Configuration policies:
  - Interfaces
  - Inline Sets
  - DHCP
  - VTEP
  - All Routing
- Decryption policy
- DNS policy
- FlexConfig
- Intrusion policy and Network Analysis Policy (NAP), Snort 3 only.
- Malware and File policy

- Network Address Translation (NAT)
- Network Discovery policy
- Platform Settings
- Prefilter
- QoS
- Umbrella SASE Topology
- VPN policies, both site-to-site and remote access
- Zero Trust Access

Supported Objects

- AAA Server
- Access List
- Address Pools
- AS Path
- Cipher suit lists
- Community List
- Distinguished Name objects
- DHCP IPv6 Pools
- DNS Server Group
- FlexConfig objects
- Group policy
- Interface
- Key Chain
- Network
- PKI certificates, all objects
- Policy List
- Port
- Prefix List
- Route Map
- Sinkhole
- SLA Monitor
- Time Range

- Time Zone
- Tunnel Zone
- URL
- Variable Set
- VLAN Tag
- VPN objects (IKEv1, IKEv2 IPSec and policy, PKI enrollment, certificate map)

Requirements and Prerequisites for Change Management

Model Support

Management Center

Supported Domains

Any

User Roles

- To enable or disable Change Management: Admin.
- To both modify and review tickets:
  - Admin
  - Network Admin
- To modify tickets only:
  - Access Admin
  - Intrusion Admin
- To review tickets only:
  - Security Approver

Guidelines and Limitations for Change Management

- When operating in change management mode, users can make changes to supported policies, but they cannot save the changes. For example, you could go through the dialog box to create a new Platform Settings policy without an open ticket, but when you click OK to actually create the policy, you will get an error and the policy will not be created.
- The following activities require that all tickets be in a terminal state, that is, approved or discarded: backup/restore, moving a device between domains, upgrading Firewall Management Center.

- Deleting a device from the inventory requires that all tickets involving that device be approved or discarded.
- Some processes, such as deployment and backup/restore, prevent you from changing the Change Management mode. Wait until the process completes to change the mode.
- Your ability to create objects while configuring a feature is constrained based on whether the feature and objects are all supported by change management. For example, importing a configuration is not supported by change management. Therefore, you cannot create security zone objects, which are supported, during the import. On the other hand, you can create new objects while configuring access control rules, because both are supported.
- When using Cloud-Delivered Firewall Management Center, a user defined in Cisco Security Cloud Control is available to be assigned tickets only after the user cross-launches cdFMC at least once. Until the first cross-launch, the user does not exist in cdFMC.

Enabling or Disabling Change Management

The default is that the Change Management workflow is disabled. Users do not need to open tickets and get approval when making configuration changes. If you want to enforce the Change Management workflow, you must enable it globally for the system.

Before you begin

There are several system processes that prevent you from enabling/disabling change management. If any of the following are in process, you need to wait for them to complete before changing these settings: backup/restore; import/export; domain movement; upgrade; Flexconfig migration; device registration; high-availability registration, creation, break, or switch; cluster create, registration, break, edit, add or remove nodes; EPM break out or join.

An access control policy cannot be locked when you change these settings. If a policy is locked, you must wait for the lock to be released before enabling/disabling this feature.

Procedure

Step 1 Choose **System** (⚙) > **Configuration**.

Step 2 Click **Change Management**.

Step 3 Select **Enable Change Management**.

To disable the feature, deselect the option. All tickets must be approved or discarded to disable Change Management. You cannot disable Change Management if any ticket is in the In Progress, On Hold, Rejected, or Pending Approval state.

Step 4 Select the **Number of approvals required**, which is how many administrators must approve the change for the ticket to be approved and deployable. The default is 1, but you can require up to 5 approvers per ticket. Users can override this number when creating tickets.

Step 5 Select the **Ticket Purge Duration**, which is the number of days to keep approved tickets, from 1-100 days. The default is 5 days.

Step 6 (Optional.) Enter the **Reply to Address** and the email addresses for the **List of Approver Addresses**. You must also configure the Email Notification system settings for email to work.

Step 7 Click **Save**.

The system adds the **Ticket** ( ) shortcut to the menu bar, and the **System** ( ) > **Change Management Workflow** command. Users can manage tickets using these methods.

Managing Tickets

When you enable Change Management, configuration changes for supported policies must be done within the context of a ticket. You open a ticket, make your changes, then submit the ticket for approval.

You can see a list of tickets, and create new ones, either on the Change Management page or through the Ticket quick-access menu. All ticket changes are synchronized in each menu, so you can switch back and forth at your convenience and use whichever method you prefer.

Note When you open a ticket and make a change to a supported policy, that policy is locked from changes by other users or through other tickets. The policy remains locked until the ticket is approved or discarded.

Procedure

Step 1 Do any of the following:

- Choose **System** ( ) > **Change Management Workflow** to open a page showing existing tickets.
- Click the **Ticket** ( ) quick-access menu. The icon can either be named Select a Ticket (if no ticket is opened), the ticket name if a ticket is opened, or unnamed if no tickets exist.

Both pages are organized the same. The **Ticket** tab lists all tickets, whereas the **Review** tab lists tickets that have been submitted for approval. The default view shows your tickets only.

Step 2 On the **Ticket** tab, take any of these actions:

- To create a new ticket, click **Add Ticket**.
- To view the details of a ticket, click the > next to the ticket name. The Details page includes UUID, name, description, user, last modified date, and comments. The History page includes the status changes for the ticket. The image at the top shows where the ticket stands in the overall workflow.
- To preview the configuration changes for an open ticket, click **Preview** ( ).
- To validate the configuration changes in an open ticket, click **Validate** ( ) or **More** ( ) > **Validate**. A dialog box opens showing error, warning, and informational messages, if there are any validation errors.
- To open a ticket, click **Open** ( ) or **More** ( ) > **Open**.

- To close an open ticket, click **Put Ticket on Hold (X)** or **More** (⋮) > **Put Ticket on Hold**. Closing a ticket does not submit it for review, nor does it release any locks placed on edited policies.
- To submit an open ticket for review and approval, click **Submit for Approval** (📝) or **More** (⋮) > **Submit for Approval**. The ticket must be open to be submitted.
- To discard a ticket, click **Discard** (🗑) or **More** (⋮) > **Discard**.
- To take over or reassign a ticket, click **More** (⋮) > **Take Over Ticket** when viewing all tickets in the system.
- To search for a ticket, type a string in the search box. The search looks at ticket name, description, and responsible user.
- To filter the list by ticket status (on the Change Management Workflow page), click the status above the list: **New**, **Open**, **On Hold** (ticket was closed), **Rejected**, **Pending Approval**, **Approved**. Each status has a count of the number of tickets in that state. Click **All** under **My Tickets** to return to the default of showing all of your tickets, or **All** under **Tickets in System** to see everyone's tickets.

Step 3 On the **Reviews** tab, take any of these actions on submitted tickets. The list is empty if there are no submitted tickets. In addition, only users with Review Ticket permissions can see this tab.

- To preview the configuration changes for a ticket, click **Preview** (🔍).
- To validate the configuration changes in an open ticket, click **Validate** (✓) or **More** (⋮) > **Validate**.
- To approve the ticket, click **Approve** (✓) or **More** (⋮) > **Approve**.
- To disapprove the ticket, click **Reject** (✗) or **More** (⋮) > **Reject**.

Creating Change Management Tickets

When using the Change Management workflow, you must do all configuration changes within the context of an open ticket. If you do not already have a ticket, you must create a new one.

Procedure

Step 1 Choose **System** (⚙) > **Change Management Workflow**, or click the **Ticket** (🎫) shortcut menu.

Step 2 Click **Add Ticket**.

Step 3 Configure the ticket options:

- **Name**—The name of the ticket. The name can include letters, numbers, spaces, and the following special characters: #-_!
- **Description**—An optional description of what you intend to configure using this ticket. For example, if you have a case number related to what you intend to fix using this ticket, that would be useful information in a description.

- **Number of Approvers**—How many administrators must approve the change for the ticket to be approved and deployable. You can specify 1-5.
- **Assign to**—Select the user who will own the ticket and be responsible for implementing the changes. Choose **self** to assign it to yourself.

Step 4 Click one of the following: .

- **Create**—The ticket is added to the list of tickets, but it is not opened. You need to open it before you can work within the context of the ticket.
- **Create and Open**—The ticket is added to the list of tickets and also opened.

Opening a Ticket for Configuration Changes

Before you can make changes within a ticket, you must open the ticket.

If you have another ticket open, the system puts it on hold (closes it) for you before opening the new one.

Procedure

Step 1 Choose **System** ( ) > **Change Management Workflow**, or click the **Ticket** ( ) shortcut menu.

Step 2 On the **Tickets** tab, click **Open** ( ) or **More** ( ) > **Open** for the ticket.

Step 3 Optionally, enter a comment for the action.

Step 4 Click **Open**.

You can now start your configuration changes. The name of the Ticket icon changes to the name of the open ticket.

Previewing a Ticket

You can preview a ticket while making your configuration changes, or before approving it. The preview shows all configuration changes made within the context of the ticket.

Procedure

Step 1 Choose **System** ( ) > **Change Management Workflow**, or click the **Ticket** ( ) shortcut menu.

Step 2 Click **Preview** ( ) for the ticket.

The Preview dialog box opens. Changes are color coded according to the legend at the top of the dialog box.

Step 3 Select the policy whose changes you want to view in the Changed Policies list.

You are shown both the current version of the policy as it is defined in Secure Firewall Management Center (on the left) and the proposed changes defined within the ticket.

For policies that contain pages, such as Platform Settings, you can select the overall policy to see all changes, or specific pages within the policy, in the Changed Policies list.

You cannot alter the changes from within the preview. If you need to change something, you must close the preview and return to the policy you want to change.

Step 4 Optionally, click **Download as PDF** to save the preview to a PDF file for offline viewing or archival purposes.

Step 5 Click **OK**.

Submitting a Ticket

After you have completed the changes needed for a ticket, you can preview and validate the changes. Then, when you are satisfied with the changes, submit the ticket for review and approval.

Changes made within the ticket are not applied until you submit the ticket and the ticket is approved. Until approval happens, all policies modified within a ticket are locked to that ticket and cannot be changed by anyone else.

Procedure

Step 1 Choose **System** ( ) > **Change Management Workflow**, or click the **Ticket** ( ) shortcut menu.

Step 2 Click **Submit for Approval** ( ) or **More** ( ) > **Submit for Approval** for the open ticket.

Step 3 Optionally, enter a comment for your action.

Step 4 Click **Submit**.

Discarding a Ticket

If you no longer need to make the changes for which you created a ticket, you can discard the ticket. When you discard a ticket, any changes that you made within the ticket are removed.

You cannot undo this action and retrieve a ticket with its changes. If you need them, you must create a new ticket and start over.

You cannot discard a ticket after you submit it. However, if the approver rejects the ticket, you can then discard it.

Note If you have permission to modify tickets, you can discard tickets that belong to another user. This makes it possible to deal with situations where an admin is on vacation or otherwise unavailable to manage an in-process ticket. If you have Review Ticket permission, you can alternatively reassign or take over the ticket rather than discard it.

Procedure

Step 1 Choose **System** ( ) > **Change Management Workflow**, or click the **Ticket** ( ) shortcut menu.

Step 2 Click **Discard** ( ) or **More** ( ) > **Discard** for the ticket.

Step 3 Optionally, enter a comment for your action.

Step 4 Click **Discard**.

Approving or Rejecting a Ticket

When a user submits a ticket, it must be approved for the changes made within the ticket to become active and available for deployment.

Whether you can approve your own tickets, or there is a separate approver, depends on your workplace policies and how user roles are assigned, not on the management software.

The Details view includes a summary of how many approvers are required for the ticket, and who has approved the ticket.

If the changes are inadequate or undesirable, you can reject the ticket. Rejected tickets go back to the submitter, who can then make additional changes and resubmit the ticket, or simply discard the ticket and the configuration changes it contains.

Procedure

Step 1 Choose **System** ( ) > **Change Management Workflow**, or click the **Ticket** ( ) shortcut menu.

Step 2 On the **Review** tab, click **Preview** ( ) for the ticket and evaluate the proposed changes.

You can also click **Validate** ( ) or **More** ( ) > **Validate** to check for errors.

Step 3 After completing your evaluation, do one of the following:

- To approve the ticket, click **Approve** ( ) or **More** ( ) > **Approve**.
- To disapprove the ticket, click **Reject** ( ) or **More** ( ) > **Reject**.

Step 4 Optionally, enter a comment for your action.

Step 5 Click **Approve** or **Reject**, as appropriate.

Taking Over or Reassigning Tickets

Sometimes it might be necessary to take over a ticket that someone else created. For example, the ticket owner might be on vacation or is otherwise unavailable, and the ticket is blocking updates that need to be deployed.

You can also use this procedure to reassign your own ticket to another person.

Before you begin

Following are the permissions required to take over tickets:

- Admin user—You can assign tickets to yourself or other users.
- Modify or Review ticket + System > User Management > Users (custom role)—You can assign tickets to yourself or other users.

However, you can assign a user only if the user has the same role as the current ticket owner, or the Admin role. This ensures that the new user has the permissions required to configure the features currently modified within the ticket.

You cannot reassign a ticket that has been submitted for approval.

Procedure

Step 1 Choose **System** (⚙) > **Change Management Workflow**.

Step 2 Click on **All** under **Tickets in System**.

Step 3 Click **More** (⋮) > **Take Over Ticket**.

Step 4 Select the user who should now own the ticket.

The list of users is restricted to those who have the permissions required to edit the policies already changed within the ticket. For example, if a ticket contains changes to the access control policy, the list of users contains only those users who are allowed to modify the access control policy.

Step 5 Enter an optional comment and click **Takeover**.

History for Change Management

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Taking over tickets and additional support for Change Management.	7.6.0	Any	You can take over another user's ticket. This is useful if a ticket is blocking other updates to a policy and the user is unavailable. In addition, the following features are now included in the Change Management approval workflow: Decryption policies, DNS policies, file and malware policies, network discovery, certificates and certificate groups, cipher suite lists, Distinguished Name objects, Sinkhole objects.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Change management.	7.4.1	Any	You can enable change management if your organization needs to implement more formal processes for configuration changes, including audit tracking and official approval before changes are deployed. We added the **System**(⚙) > **Configuration** > **Change Management** page to enable the feature. When enabled, there is a **System**(⚙) > **Change Management Workflow** page, and a new **Ticket**(📋) quick access icon in the menu.

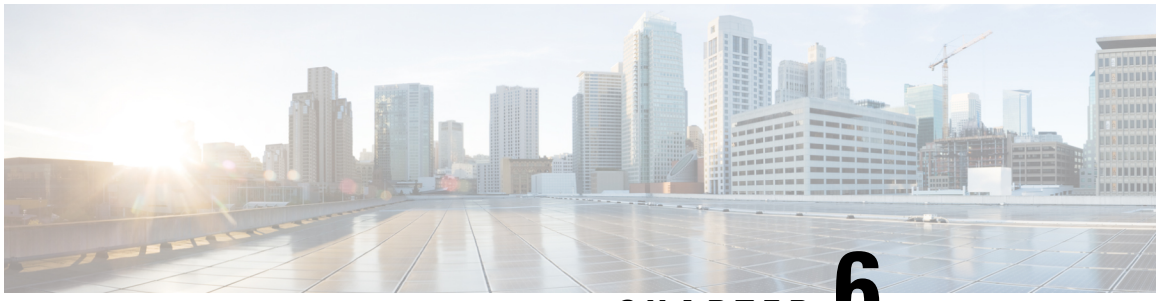

CHAPTER 6

Configuration Deployment

This chapter describes how to download configuration changes to one or more managed devices.

- About Configuration Deployment, on page 247
- Requirements and Prerequisites for Policy Management, on page 257
- Best Practices for Deploying Configuration Changes, on page 258
- Deploy the Configuration, on page 259
- Manage Deployments, on page 266
- History for Configuration Deployment, on page 279

About Configuration Deployment

All device configuration is managed by the Firewall Management Center and then deployed to the managed devices.

Configuration Changes that Require Deployment

The system marks out-of-date policies with red status text that indicates how many of its targeted devices need a policy update. To clear this status, you must re-deploy the policy to the devices.

Deployment Required

Configuration changes that require a deployment include:

- Modifying an access control policy: any changes to access control rules, the default action, policy targets, Security Intelligence filtering, advanced options including preprocessing, and so on.
- Modifying any of the policies that the access control policy invokes: the SSL policy, network analysis policies, intrusion policies, file policies, identity policies, or DNS policies.
- Changing any reusable object or configuration used in an access control policy or policies it invokes:
  - network, port, VLAN tag, URL, and geolocation objects
  - Security Intelligence lists and feeds
  - application filters or detectors
  - intrusion policy variable sets

    - file lists
    - decryption-related objects and security zones
- Updating the system software, intrusion rules, or the vulnerability database (VDB).

Keep in mind that you can change some of these configurations from multiple places in the web interface. For example, you can modify security zones using the object manager (**Objects** > **Object Management**), but modifying an interface type in a device's configuration (**Devices** > **Device Management**) can also change a zone and require a deployment.

Deployment Not Required

Note that the following updates do **not** require a deployment:

- automatic updates to Security Intelligence feeds and additions to the Security Intelligence global Block or Do Not Block list using the context menu
- automatic updates to URL filtering data
- scheduled geolocation database (GeoDB) updates

Deployment Preview

Preview provides a snapshot of all the policy and object changes to be deployed on the device. The policy changes include the new policies, changes in the existing policies, and the deleted policies. The object changes include the added and modified objects which are used in policies. The unused object changes are not displayed because they are not deployed on the device.

Click the **Preview** icon next to a deployment job to view the configuration change log for that are pending to be deployed on the device. The change log includes:

- **Comparison View**: Displays a side-by-side comparison of all the device configuration changes made since your last deployment.
- **Advanced View**: Displays the pending CLI commands to be applied to the device.

For more information about viewing the deployment preview, see Deploy Configuration Changes, on page 259.

The preview shows all the default values, even when they are not altered, along with the other configured settings when an interface or a platform settings policy is added for the first time. Similarly, the high availability-related policies and default values for settings are shown, even when they are not altered, in the first preview after a high availability pair is configured or disrupted.

To view changes due to an auto rollback, see Edit Deployment Settings, on page 201.

Unsupported Features

- Object additions and attribute changes are displayed in the preview only if the objects are associated with any device or interface. Object deletions are not displayed.
- Preview is not supported for the following policies:
    - High availability

- Network discovery
- Network analysis
- Device settings

- User information at the rule level is not available for intrusion policies.
- The preview does not show the reordering of rules across policies.

  For DNS policies, reordered rules appear in the preview list as rule additions and deletions. For example, moving a rule from position 1 to position 3 in the rule order is displayed as if the rule was deleted from position 1 and added as a new rule in position 3. Similarly, when a rule is deleted, the rules under it are listed as edited rules as they have changed their positions. The changes are displayed in the final order in which they appear in the policy.
- Preview is not supported in the following HA scenarios:
  - If a device was in standalone mode and if a chain is made, then an auto-deployment is triggered. For that particular job, preview is not supported. On hover over the **Preview** (), a message is displayed that it is a HA bootstrap deployment, and no preview is supported.
  - **Configuration groups** - Consider a flow in which a device was initially standalone. Subsequently, three deployments took place. In the fourth deployment, the device was a HA bootstrap deployment. After these, the user deploys devices 5, 6, and 7. The deployment 7 is an HA break deployment, and the user deploys devices 8, 9, and 10.

    In this flow, the preview between 3 and 5 is not supported because 4 was a HA deployment. Similarly, the preview between 8 and 3 is also not supported. Preview is supported only from 3 to 1, 7,6, 5, 4, and 10, 9, and 8.
  - If a device is broken (HA is broken) then the new device is considered as a fresh device.

Selective Policy Deployment

The Firewall Management Center allows you to select a specific policy within the list of all the changes on the device that are due for deployment and deploy only the selected policy. Selective deployment is available only for the following policies:

- Access control policies
- Intrusion policies
- Malware and file policies
- DNS policies
- Identity policies
- SSL policies
- QoS policies
- Prefilter policies
- Network discovery

- NAT policies
- Routing policies
- VPN policies

There are certain limitations to selectively deploying policies. Follow the contents in the table below to understand when selective policy deployment can be used.

Table 11: Limitations for Selective Deployment

Type	Description	Scenarios
Full deployment	Full deployment is necessary for specific deploy scenarios, and the Firewall Management Center does not support selective deployment in such scenarios. If you encounter an error in such scenarios, you may choose to proceed by selecting all the changes for deployment on the device.	Scenarios wherein a full deployment is required are: • The first deployment after you have upgraded the Firewall Threat Defense or the Firewall Management Center. • The first deployment after you have restored the Firewall Threat Defense. • The first deployment after modifications in the Firewall Threat Defense interface settings. • The first deployment after modifications in the virtual router settings. • When the Firewall Threat Defense device is moved to a new domain (global to sub-domain or sub-domain to global).
Associated policy deployment	The Firewall Management Center identifies interdependent policies which are interlinked. When one of the interlinked policies is selected, the remaining interlinked policies are automatically selected.	Scenarios wherein an associated policy is automatically selected: • When a new object is associated with an existing policy. • When an existing policy's object is modified. Scenarios wherein multiple policies are automatically selected: • When a new object is associated with an existing policy, and the same object is already associated with other policies, all the associated policies are automatically selected. • When a shared object is modified, all the associated policies are automatically selected.

Type	Description	Scenarios
Interdependent policy changes (shown using color-coded tags)	The Firewall Management Center dynamically detects dependencies in-between policies, and between the shared objects and the policies. The interdependency of the objects or policies is shown using color-coded tags.	Scenarios wherein color-coded interdependent policies or objects are automatically selected: • When all the out-of-date policies have interdependent changes. For example, when an access control policy, an intrusion policy, and a NAT policy are out-of-date. Since access control policy and NAT policy share an object, all policies are selected together for deployment. • When all out-of-date policies share an object, and the object is modified.
Access Policy Group specifications	Access Policy Group policies are listed together in the preview window under **Access Policy Group** when you click **View** ( ).	The scenarios and the expected behavior for Access Policy Group policies are: • If the access control policy is out-of-date, all other out-of-date policies under this group, except file policy and intrusion policy, are selected when the access control policy is selected for deployment. However, if the access control policy is out-of-date, intrusion and file policies can be individually selected or deselected irrespective of whether the access control policy is selected or not, unless there are any dependent changes. For example, if a new intrusion policy is assigned to an access control rule, it indicates that there are dependent changes, then both the access control policy and the intrusion policy will be automatically selected when either of them is selected. • If no access control policy is out-of-date, other out-of-date policies in this group can be selected and deployed individually.

System Username

The Firewall Management Center displays the username as **system** for the following operations:

- Rollback
- Upgrade
- Firewall Threat Defense backup and restore
- SRU update

- LSP update
- VDB update

Auto-Enabling of Application Detectors

If you are performing application control but disable required detectors, the system automatically enables the appropriate system-provided detectors upon policy deploy. If none exist, the system enables the most recently modified user-defined detector for the application.

Asset Rediscovery with Network Discovery Policy Changes

When you deploy changes to a network discovery policy, the system deletes and then rediscovers MAC address, TTL, and hops information from the network map for the hosts in your monitored networks. Also, the affected managed devices discard any discovery data that has not yet been sent to the Firewall Management Center.

Snort Restart Scenarios

When the traffic inspection engine referred to as *the Snort process* on a managed device restarts, inspection is interrupted until the process resumes. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information. Additionally, resource demands may result in a small number of packets dropping without inspection when you deploy, regardless of whether the Snort process restarts.

Any of the scenarios in the following table cause the Snort process to restart.

Table 12: Snort Restart Scenarios

Restart Scenario	More Information
Deploying a specific configuration that requires the Snort process to restart.	Configurations that Restart the Snort Process When Deployed or Activated, on page 256
Modifying a configuration that immediately restarts the Snort process.	Changes that Immediately Restart the Snort Process, on page 257
Traffic-activation of the currently deployed Automatic Application Bypass (AAB) configuration.	Configure Automatic Application Bypass, on page 198
Enabling or disabling "Logging connection events to RAM disk" feature.	See the section **Log to Ramdisk** in Troubleshoot Drain of FMC Unprocessed Events.

Related Topics

Access Control Policy Advanced Settings, on page 1756
Configurations that Restart the Snort Process When Deployed or Activated, on page 256

Restart Warnings for Devices

When you deploy, the **Inspect Interruption** column in the deploy page specifies whether a deployed configuration restarts the Snort process on the Firewall Threat Defense device. When the traffic inspection

engine referred to as *the Snort process* restarts, inspection is interrupted until the process resumes. Whether traffic is interrupted or passes without inspection during the interruption depends on how the device handles traffic. Note that you can proceed with the deployment, cancel the deployment and modify the configuration, or delay the deployment until a time when deploying would have the least impact on your network.

When the **Inspect Interruption** column indicates **Yes** and you expand the device configuration listing, the system indicates any specific configuration type that would restart the Snort process with an **Inspect Interruption** (). When you hover your mouse over the icon, a message informs you that deploying the configuration may interrupt traffic.

The following table summarizes how the deploy page displays inspection interruption warnings.

Table 13: Inspection Interruption Indicators

Type	Inspect Interruption	Description
Firewall Threat Defense	**Inspect Interruption** ()Yes	At least one configuration would interrupt inspection on the device if deployed, and might interrupt traffic depending on how the device handles traffic. You can expand the device configuration listing for more information.
	--	Deployed configurations will not interrupt traffic on the device.
	Undetermined	The system cannot determine if a deployed configuration may interrupt traffic on the device. Undetermined status is displayed before the first deployment after a software upgrade, or in some cases during a Support call.
	Errors ()	The system cannot determine the status due to an internal error. Cancel the operation and click **Deploy** again to allow the system to redetermine the **Inspect Interruption** status. If the problem persists, contact Support.
sensor	--	The device identified as *sensor* is not the Firewall Threat Defense device; the system does not determine if a deployed configuration may interrupt traffic on this device.

For information on all configurations that restart the Snort process for all device types, see Configurations that Restart the Snort Process When Deployed or Activated, on page 256.

Inspect Traffic During Policy Apply

Inspect traffic during policy apply is an advanced access control policy general setting that allows managed devices to inspect traffic while deploying configuration changes; this is the case unless a configuration that you deploy requires the Snort process to restart. You can configure this option as follows:

- Enabled — Traffic is inspected during the deployment unless certain configurations require the Snort process to restart.

  When the configurations you deploy do not require a Snort restart, the system initially uses the currently deployed access control policy to inspect traffic, and switches during deployment to the access control policy you are deploying.

- Disabled — Traffic is not inspected during the deployment. The Snort process always restarts when you deploy.

The following graphic illustrates how Snort restarts can occur when you enable or disable **Inspect traffic during policy apply**.

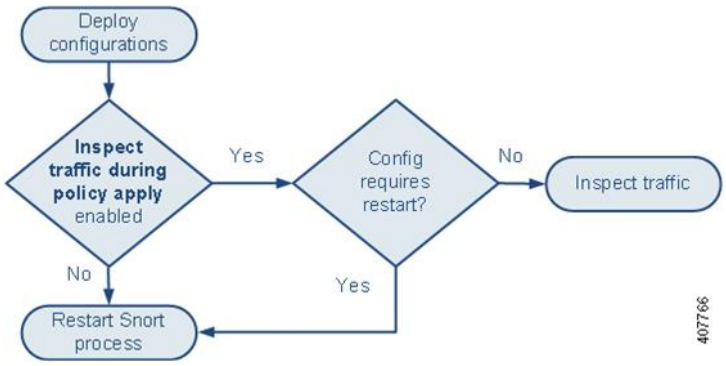

Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort Restart Traffic Behavior, on page 254 and Configurations that Restart the Snort Process When Deployed or Activated, on page 256.

Snort Restart Traffic Behavior

The following tables explain how different devices handle traffic when the Snort process restarts.

Table 14: The Firewall Threat Defense and the Firewall Threat Defense Virtual Restart Traffic Effects

Interface Configuration	Restart Traffic Behavior
inline: **Snort Fail Open: Down**: disabled	dropped
inline: **Snort Fail Open: Down**: enabled	passed without inspection Some packets can be delayed in buffer for several seconds before the system recognizes that Snort is down. This delay can vary depending upon the load distribution. However, the buffered packets are eventually passed.

Interface Configuration	Restart Traffic Behavior
routed, transparent (including EtherChannel, redundant, subinterface): **preserve-connection** enabled (**configure snort preserve-connection enable**; default) For more information, see Cisco Secure Firewall Threat Defense Command Reference.	existing TCP/UDP flows: passed without inspection so long as at least one packet arrives while Snort is down new TCP/UDP flows and all non-TCP/UDP flows: dropped Note that the following traffic drops even when **preserve-connection** is enabled: • plaintext, passthrough prefilter tunnel traffic that matches an **Analyze** rule action or an **Analyze all tunnel traffic** default policy action • connections that do not match an access control rule and are instead handled by the default action. • decrypted TLS/SSL traffic • a safe search flow • a captive portal flow
routed, transparent (including EtherChannel, redundant, subinterface): **preserve-connection** disabled (**configure snort preserve-connection disable**)	dropped
inline: tap mode	egress packet immediately, copy bypasses Snort
passive	uninterrupted, not inspected

Note In addition to traffic handling when the Snort process is down while it restarts, traffic can also pass without inspection or drop when the Snort process is busy, depending on the configuration of the Snort Fail Open **Busy** option (see Configure an Inline Set, on page 901). A device supports either the Failsafe option or the Snort Fail Open option, but not both.

Note When the Snort process is busy but not down during configuration deployment, some packets may drop on routed, switched, or transparent interfaces if the total CPU load exceeds 60 percent.

Warning Do not reboot the system while the Snort Rule Update is in progress.

Snort-busy drops happen when snort is not able to process the packets fast enough. Lina does not know whether Snort is busy due to processing delay, or if is stuck or due to call blocking. When transmission queue is full, snort-busy drops occur. Based on Transmission queue utilization, Lina will try to access if the queue is being serviced smoothly.

Configurations that Restart the Snort Process When Deployed or Activated

Deploying any of the following configurations except AAB restarts the Snort process as described. Deploying AAB does not cause a restart, but excessive packet latency activates the currently deployed AAB configuration, causing a partial restart of the Snort process.

Access Control Policy Advanced Settings

- Deploy when **Inspect Traffic During Policy Apply** is disabled.
- Add or remove an SSL policy.

File Policy

Deploy the first or last of any one of the following configurations; note that while otherwise deploying these file policy configurations does not cause a restart, deploying non-file-policy configurations can cause restarts.

- Take either of the following actions:
  - Enable or disable **Inspect Archives** when the deployed access control policy includes at least one file policy.
  - Add the first or remove the last file policy rule when **Inspect Archives** is enabled (note that at least one rule is required for **Inspect Archives** to be meaningful).
- Enable or disable **Store files** in a **Detect Files** or **Block Files** rule.
- Add the first or remove the last active file rule that combines the **Malware Cloud Lookup** or **Block Malware** rule action with an analysis option (**Spero Analysis or MSEXE**, **Dynamic Analysis**, or **Local Malware Analysis**) or a store files option (**Malware, Unknown**, **Clean**, or **Custom**).

Note that access control rules that deploy these file policy configurations to security zones or tunnel zones cause a restart only when your configuration meets the following conditions:

- Source or destination security zones in your access control rule must match the security zones associated with interfaces on the target devices.
- Unless the destination zone in you access control rule is *any*, a source tunnel zone in the rule must match a tunnel zone assigned to a tunnel rule in the prefilter policy.

Identity Policy

- When SSL decryption is disabled (that is, when the access control policy does not include an SSL policy), add the first or remove the last active authentication rule.

  An active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive or VPN identity cannot be established** selected.

Network Discovery

- Enable or disable non-authoritative, traffic-based user detection over the HTTP, FTP, or MDNS protocols, using the network discovery policy.

Device Management

- MTU: Change the highest MTU value among all non-management interfaces on a device.
- Automatic Application Bypass (AAB): The currently deployed AAB configuration activates when a malfunction of the Snort process or a device misconfiguration causes a single packet to use an excessive amount of processing time. The result is a partial restart of the Snort process to alleviate extremely high latency or prevent a complete traffic stall. This partial restart causes a few packets to pass without inspection, or drop, depending on how the device handles traffic.

Updates

- System update: Deploy configurations the first time after a software update that includes a new version of the Snort binary or data acquisition library (DAQ).
- For managed devices running Snort 3, deploying configurations the first time after installing a vulnerability database (VDB) update may temporarily interrupt application detection, but there will be no traffic interruptions.

Related Topics

Deploy Configuration Changes, on page 259
Snort Restart Scenarios, on page 252

Changes that Immediately Restart the Snort Process

The following changes immediately restart the Snort process without going through the deploy process. How the restart affects traffic depends on how the target device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

- Take any of the following actions involving applications or application detectors:
  - Activate or deactivate a system or custom application detector.
  - Delete an activated custom detector.
  - **Save and Reactivate** an activated custom detector.
  - Create a user-defined application.

  A message warns you that continuing restarts the Snort process, and allows you to cancel; the restart occurs on any managed device in the current domain or in any of its child domains.
- Create or break a Firewall Threat Defense high availability pair.

  A message warns you that continuing to create a high availability pair restarts the Snort process on the primary and secondary devices and allows you to cancel.

Requirements and Prerequisites for Policy Management

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Network Admin
- Security Approver

Best Practices for Deploying Configuration Changes

The following are guidelines for deploying configuration changes.

Reliable Management Connection

The management connection between the Firewall Management Center and the device is a secure, TLS-1.3-encrypted communication channel between itself and the device.

You do not need to run this traffic over an additional encrypted tunnel such as Site-to-Site VPN for security purposes. If the VPN goes down, for example, you will lose your management connection, so we recommend a simple management path.

Caution We recommend against a device's management connection going through a VPN tunnel that terminates on the device itself. If you deploy a configuration change that causes the VPN to go down, the management connection will be disconnected and you will not have any way to recover the configuration without connecting directly to the device.

If management traffic exits a VPN-terminating interface, be sure to exclude the management traffic from the VPN tunnel.

Maximum Concurrent Deployments

You should not deploy to more than 25% of the maximum devices allowed for a Firewall Management Center in the same job. For example, for the FMCv300, the maximum job size should be 75 devices (25% of 300). Concurrent deployment to more devices can cause performance issues.

Deployment of Shared Policies

For best performance, deploy to devices that use the same policies. Create separate deployment jobs for each group of devices that share policies.

Time to Deploy and Memory Limitations

The time it takes to deploy depends on multiple factors, including (but not limited to):

- The configurations you send to the device. For example, if you dramatically increase the number of Security Intelligence entries you block, deployment can take longer.
- Device model and memory. On lower-memory devices, deploying can take longer.

Do not exceed the capability of your devices. If you exceed the maximum number of rules or policies supported by a target device, the system displays a warning. The maximum depends on a number of factors—not only memory and the number of processors on the device, but also on policy and rule complexity. For information on optimizing policies and rules, see Best Practices for Access Control Rules, on page 1738.

Use a Maintenance Window to Lessen the Impact of Traffic Interruptions

We *strongly* recommend you deploy in a maintenance window or at a time when interruptions will have the least impact.

When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort Restart Traffic Behavior, on page 254 and Configurations that Restart the Snort Process When Deployed or Activated, on page 256.

For the Firewall Threat Defense devices, the **Inspect Interruption** column in the Deploy dialog warns you when deploying might interrupt traffic flow or inspection. You can either proceed with, cancel, or delay deployment; see Restart Warnings for Devices, on page 252 for more information.

Related Topics

Snort Restart Scenarios, on page 252

Deploy the Configuration

After you configure your deployment, and any time you change that configuration, you must deploy the changes to affected devices. You can view deployment status in the Message Center.

Deploying updates the following components:

- Device and interface configurations
- Device-related policies: NAT, VPN, QoS, platform settings
- Access control and related policies: DNS, file, identity, intrusion, network analysis, prefilter, SSL
- Network discovery policy
- Intrusion rule updates
- Configurations and objects associated with any of these elements

You can configure the system to deploy automatically by scheduling a deploy task or by setting the system to deploy when importing intrusion rule updates. Automating policy deployment is especially useful if you allow intrusion rule updates to modify system-provided base policies for intrusion and network analysis. Intrusion rule updates can also modify default values for the advanced preprocessing and performance options in your access control policies.

Deploy Configuration Changes

After you change configurations, deploy them to the affected devices. We *strongly* recommend that you deploy in a maintenance window or at a time when any interruptions to traffic flow and inspection will have the least impact.

Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort Restart Traffic Behavior, on page 254 and Configurations that Restart the Snort Process When Deployed or Activated, on page 256.

Before you begin

- Be sure all managed devices use the same revision of the Security Zones object. If you have edited security zone objects: Do not deploy configuration changes to any device until you edit the zone setting for interfaces on *all* devices you want to sync. You must deploy to all managed devices at the same time.
- To preview the deployment changes, enable REST API access. To enable the REST API access, follow the steps in *Enabling REST API Access* in the Cisco Secure Firewall Management Center Administration Guide.

Note The deployment process fails if the device configuration is being read at the device CLI during deployment. Do not execute commands such as **show running-config** during the deployment.

Procedure

Step 1 On the Firewall Management Center menu bar, click **Deploy**.

Step 2 For a quick deployment, check specific devices and then click **Deploy**, or click **Deploy All** to deploy to all devices. Otherwise, for additional deployment options, click **Advanced Deploy**.

The rest of this procedure applies to the **Advanced Deploy** screen.

Figure 97: Quick Deploy

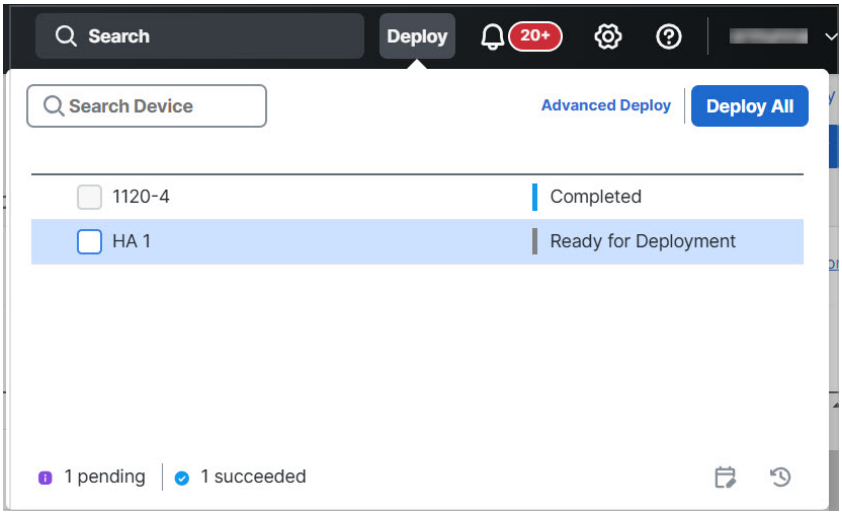

Figure 98: Advanced Deploy

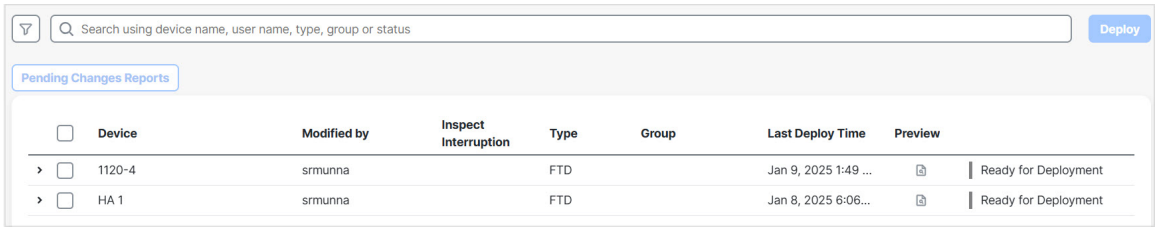

Step 3 Click **Expand Arrow** (›) to view device-specific configuration changes to be deployed.

Figure 99: Expand

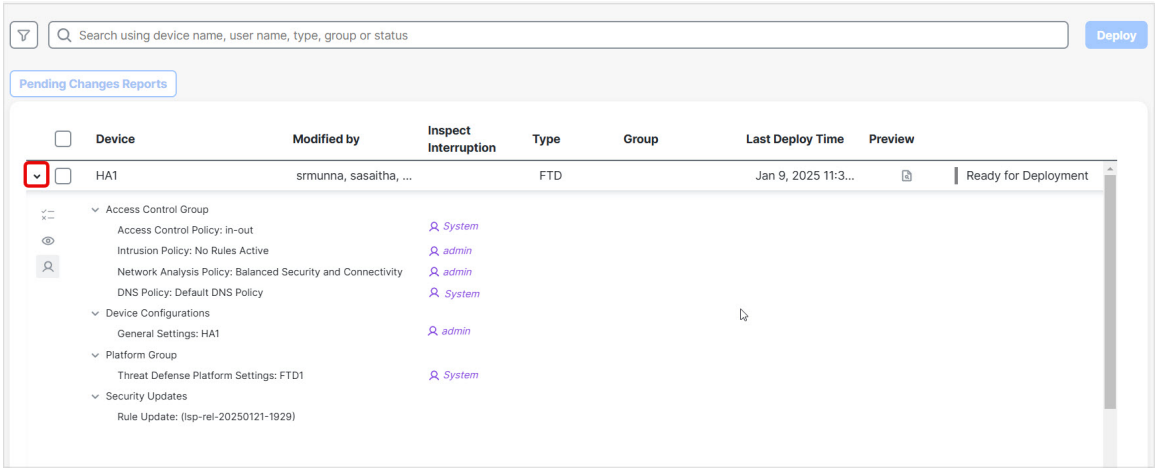

- The **Modified By** column lists the users who have modified the policies or objects. On expanding the device listing, you can view the users who have modified the policies against each policy listing. For information about when the **System** user is shown (instead of the logged-in user), see System Username, on page 251.

  Note
  Usernames are not provided for deleted policies and objects.

- The **Inspect Interruption** column indicates if traffic inspection interruption may be caused in the device during deployment.

  When the status indicates (Yes) that deploying will interrupt inspection, and perhaps traffic, on the Firewall Threat Defense device, the expanded list indicates the specific configurations causing the interruption with the **Inspect Interruption** (⚡).

  If the entry is blank in this column for a device, then it indicates that there will be no traffic inspection interruptions on that device during deployment.

  See Restart Warnings for Devices, on page 252 for information to help you identify configurations that interrupt traffic inspection and might interrupt traffic when deployed to the Firewall Threat Defense devices.

- The **Last Modified Time** column specifies when you last made the configuration changes.
- The **Preview** column allows you to preview the changes for the next deployment.

- The **Status** column provides the status for each deployment. For more information, see View Deployment Status, on page 266.

Step 4 In the **Preview** column, click **Preview** ( ) to see the configuration changes that you can deploy.

Figure 100: Preview

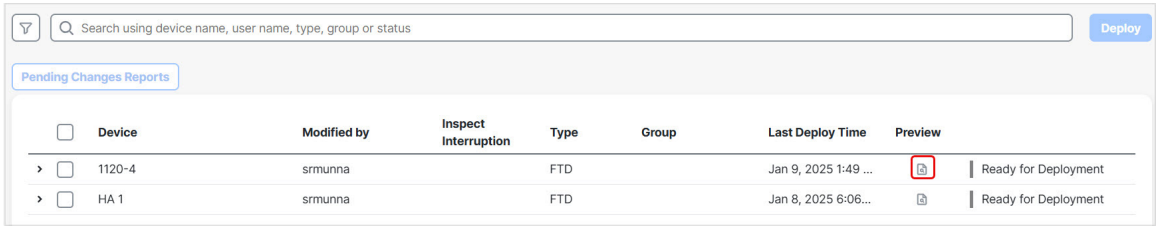

Note
If you change the Firewall Management Center name in **System** ( ) > **Configuration** > **Information**, the deployment preview does not specify this change, yet it requires a deployment.

For unsupported features for Preview, see Deployment Preview, on page 248.

The **Comparison View** tab lists all the policy and object changes. The left pane lists all the different policy types that have changed on the device, organized in a tree structure.

Figure 101: Comparison View

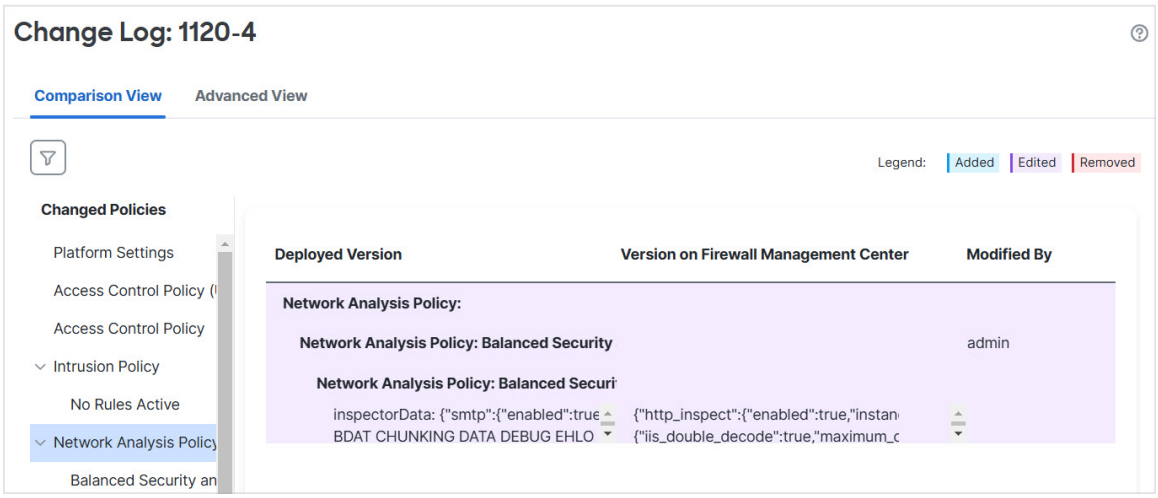

The **Filter**( ) lets you filter the policies at the user level and policy level.

The right pane lists all the additions, changes, or deletions in the policy, or the object selected in the left pane. The two columns on the right pane provide the last deployed configuration settings (in the **Deployed Version** column) versus the changes that are due for deployment (in the **Version on Firewall Management Center** column). The last-deployed configuration settings are derived from a snapshot of the last saved deployment in the Firewall Management Center and not from the device. The background colors of the settings are color-coded as per the legend available on the top-right of the page.

The **Modified By** column lists the users who have modified, or added the configuration settings. At the policy level, the Firewall Management Center displays all the users who have modified the policy, and at the rule level, the Firewall Management Center displays only the last user who has modified the rule.

You can download a copy of the change log by clicking the **Download Report** button.

The **Advanced View** tab shows the CLI commands that will be applied. This view is useful if you are familiar with ASA CLI, which is used on the back end of the Firewall Threat Defense.

Figure 102: Advanced View

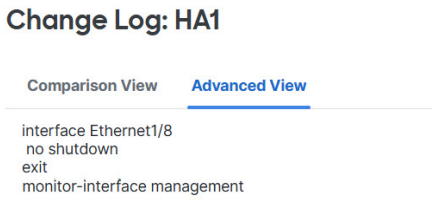

Step 5 Use **Show or Hide Policy** ( ) to selectively view or hide the associated unmodified policies.

Figure 103: Show or Hide Policy

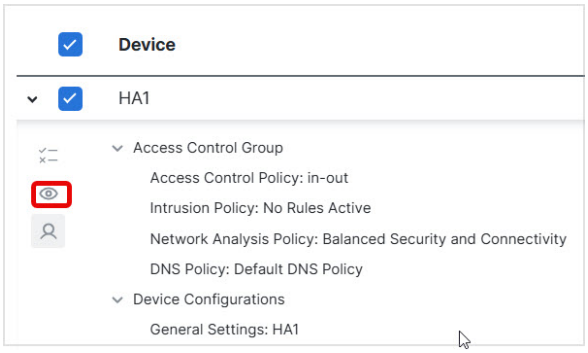

Step 6 Check the box next to the device name to deploy all configuration changes, or click **Policy selection** ( ) to select individual policies or configurations to deploy while withholding the remaining changes without deploying them.

You can also view the interdependent changes for a certain policy or configuration using this option. The Firewall Management Center dynamically detects dependencies between policies (for example, between an access control policy and an intrusion policy), and between the shared objects and the policies. Interdependent changes are indicated using color-coded tags to identify a set of interdependent deployment changes. When one of the deployment changes is selected, the interdependent changes are automatically selected.

For more details, see Selective Policy Deployment, on page 249.

Note

- When the changes in shared objects are deployed, the impacted policies should also be deployed along with them. When you select a shared object during deployment, the impacted policies are automatically selected.
- Selective deployment is not supported for scheduled deployments and deployments using REST APIs. You can only opt for complete deployment of all the changes in these cases.
- The pre-deployment checks for warnings and errors are performed not only on the selected policies, but on all the policies that are out-of-date. Therefore, the warnings or errors list shows the deselected policies as well.

- Similarly, the **Inspect Interruption** column indication on the Deployment page considers all out-of-date policies and not just the selected policies. For information on the **Inspect Interruption** column, see Restart Warnings for Devices, on page 252.

Step 7 After you select the devices or policies to deploy, click **Estimate** to get a rough estimate of the deployment duration.

Figure 104: Estimate

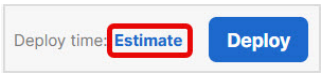

Figure 105: Deploy Time

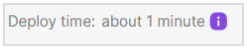

The time duration is a rough estimate (having around 70% accuracy), and the actual time taken for deployment may vary for a few scenarios. The estimate is dependable for deployments of up to 20 devices.

When an estimate is not available, it indicates that the data is not available, since the first successful deployment on the selected device is pending. This situation could occur after the Firewall Management Center reimage, version upgrade, or after a high availability failover.

Note
The estimate is incorrect and unreliable for bulk policy changes (in case of bulk policy migrations), and selective deployments because the estimate is based on the heuristic technique.

Step 8 Click **Deploy**.

Step 9 If the system identifies errors or warnings in the changes to be deployed, it displays them in the **Validation Messages** window. To view complete details, click the arrow icon before the warnings or errors.

You have the following choices:

- Deploy—Continue deploying without resolving warning conditions. Check the **Ignore warnings** checkbox, to ignore warnings and deploy the changes. You cannot proceed if the system identifies errors.
- Close—Exit without deploying. Resolve the error and warning conditions, and attempt to deploy the configuration again.

What to do next

- (Optional) Monitor deployment status; see *Viewing Deployment Messages* in the Cisco Secure Firewall Management Center Administration Guide.
- If the deployment fails, see Best Practices for Deploying Configuration Changes, on page 258.
- During deployment, if there are specific configuration changes in the deployment, the deployment failure may lead to traffic being interrupted. For example, in a cluster environment, an erroneous configuration of an IP address that is not in the same subnet as the Site IPs is configured on the interface. Due to this error, deployment fails and the device attempts to clear the configuration while the rollback operation is being processed. These events collectively lead to a deployment failure that interrupts the traffic.

See the following table to know what configuration changes may cause traffic interruption when deployment fails.

Configuration Changes	Exists?	Traffic Impacted?
Threat Defense Service changes in an access control policy	Yes	Yes
VRF	Yes	Yes
Interface	Yes	Yes
QoS	Yes	Yes

Note The configuration changes interrupting traffic during deployment is valid only if both the Firewall Management Center and the Firewall Threat Defense are of version 6.2.3 or higher.

Related Topics

Snort Restart Scenarios, on page 252

Redeploy Existing Configurations to a Device

You can force-deploy existing (unchanged) configurations to a single managed device. We *strongly* recommend you deploy in a maintenance window or at a time when any interruptions to traffic flow and inspection will have the least impact.

Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort Restart Traffic Behavior, on page 254 and Configurations that Restart the Snort Process When Deployed or Activated, on page 256.

Before you begin

Review the guidelines described in Best Practices for Deploying Configuration Changes, on page 258.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **Edit** (✎) next to the device where you want to force deployment.

Step 3 Click **Device**.

Step 4 Click **Edit** (✎) next to the **General** section heading.

Step 5 Click **Force Deploy** (→).

Note
Force-deploy takes more time than the regular deployment because it involves the complete generation of the policy rules to be deployed on the Firewall Threat Defense.

Step 6 Click **Deploy**.

The system identifies any errors or warnings with the configurations you are deploying. You can click **Proceed** to continue without resolving warning conditions. However, you cannot proceed if the system identifies an error.

What to do next

- (Optional) Monitor deployment status; see *Viewing Deployment Messages* in the Cisco Secure Firewall Management Center Administration Guide.
- If deploy fails, see Best Practices for Deploying Configuration Changes, on page 258.

Related Topics

Snort Restart Scenarios, on page 252

Manage Deployments

View Deployment Status

On the Deployment page, the **Status** column provides the deployment status for each device. If a deployment is in progress, then the live status of the deployment progress is displayed, else one of the following statuses is displayed:

- Pending—Indicates that there are changes in the device that are to be deployed.
- Warnings or errors—Indicates that the pre-deployment checks have identified warnings or errors for the deployment, and you have not proceeded with the deployment. You can continue with the deployment if there are any warnings, but not if there are any errors.

Note The status column provides the warning or error status only for a single user session on the deployment page. If you navigate away from the page or refresh the page, the status changes to pending.

- Failed—Indicates that the previous deployment attempt failed. Click on the status to view the details.
- In queue—Indicates that deployment is initiated, and the system is yet to start the deployment process.
- Completed—Indicates that deployment has completed successfully.

View Deployment History

In the deployment history, the last 10 successful deployments, the last 5 failed deployments, and last 5 rollback deployments are captured.

Procedure

Step 1 On the Firewall Management Center menu bar, click **Deploy** and then click **Deployment History** ( ).

Figure 106: Deployment History Icon

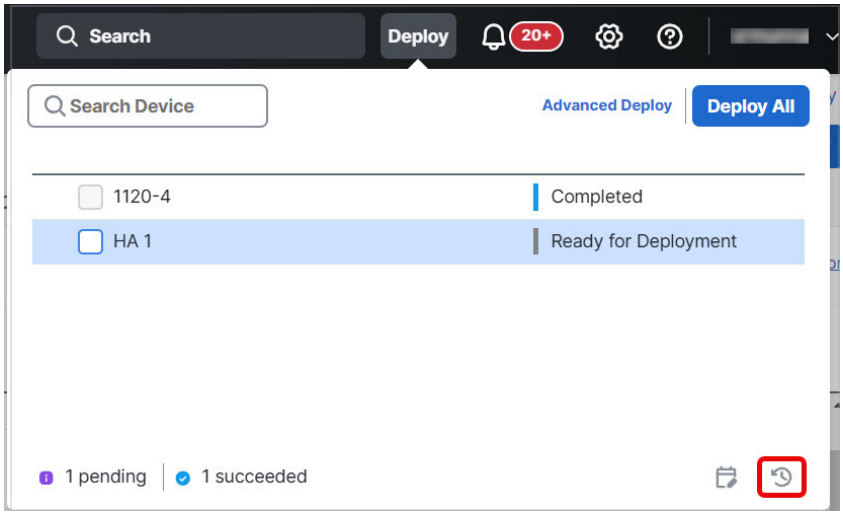

A list of all the previous deployment and rollback jobs is displayed in reverse chronological order.

Figure 107: Deployment History Page

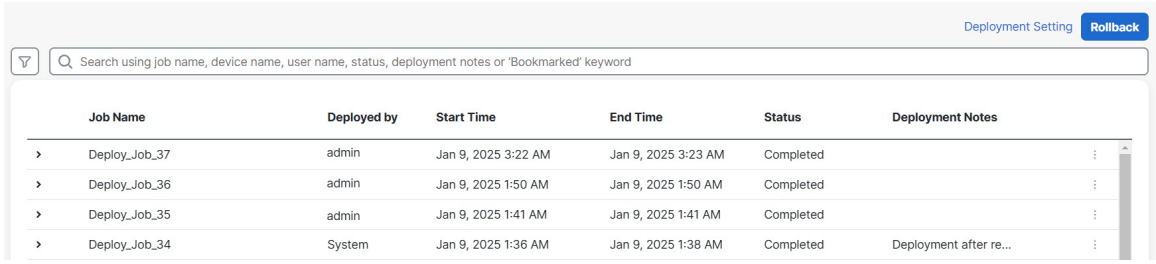

Step 2 Click **Expand Arrow** (>) next to the required deployment job to view the devices included in the job and their deployment statuses.

Figure 108: Expand

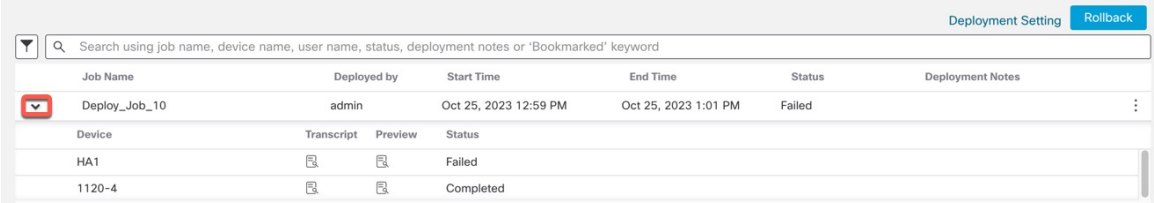

- View notes in the **Deployment Notes** column.

  Deployment notes are custom notes that a user can add as part of the deployment, and these notes are optional.

Step 3 (Optional) Click **Transcript Details** ( ) to view the commands sent to the device, and the responses received.

Figure 109: Transcript Details Icon

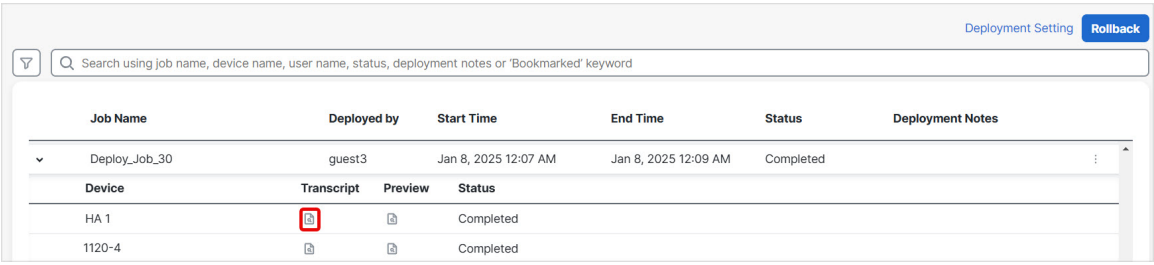

Figure 110: Transcript Details

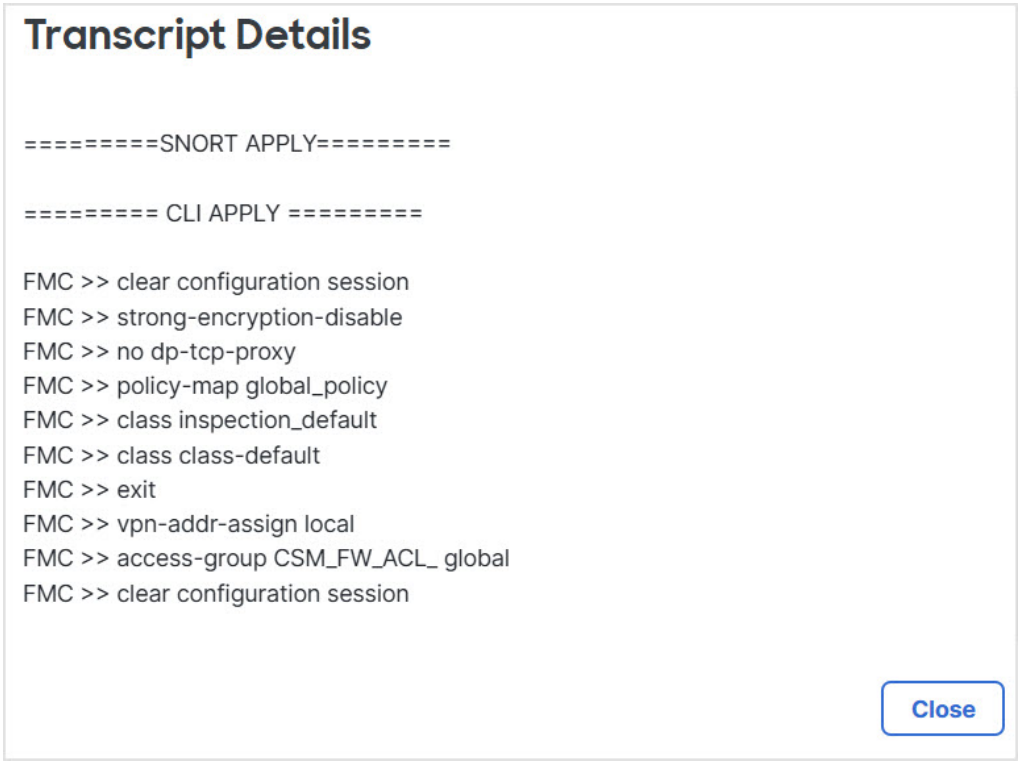

The transcript includes the following sections:

- **Snort Apply**—If there are any failures or responses from Snort-related policies, then the messages are displayed in this section. Normally, the section is empty.
- **CLI Apply**—This section covers features that are configured using commands that are sent to the device.

  Note
  The transcript for the rollback operation does not provide the CLI commands information. To view the rollback commands, see View the Deployment Rollback Transcript, on page 276.

- **Infrastructure Messages**—This section shows the status of different deployment modules.

In the **CLI Apply** section, the deployment transcript includes commands that are sent to the device, and any responses returned from the device. These responses can be informative messages or error messages. For failed deployments, look for messages that indicate errors with the commands. Examining these errors can be particularly helpful if you are using FlexConfig policies to configure customized features. These errors can help you correct the script in the FlexConfig object that is trying to configure the commands.

Note
There is no distinction that is made in the transcript between commands that are sent for managed features and those generated from FlexConfig policies.

For example, the following sequence shows that Firewall Management Center sent commands to configure GigabitEthernet0/0 with the logical name **outside**. The device responded that it automatically set the security level to 0. Firewall Threat Defense does not use the security level for anything.

```
========= CLI APPLY =========

FMC >> interface GigabitEthernet0/0
FMC >>  nameif outside
FTDv 192.168.0.152 >> [info] : INFO: Security level for "outside" set to 0 by default.
```

Step 4 (Optional) Click **Preview** ( ) to view the policy and object changes deployed on the device versus the previously deployed version.

Figure 111: Preview Icon

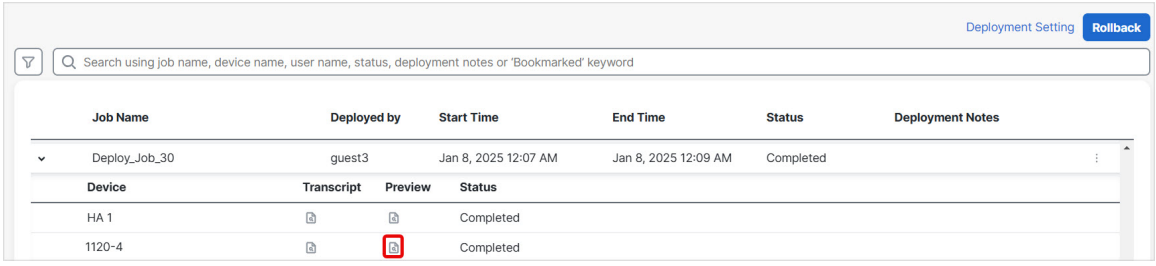

a. To compare any two versions and view the change log, choose the required versions in the drop-down boxes and click the **Show** button. The drop-down boxes show the deployment job name and the end time of the deployment.

Figure 112: Compare Versions

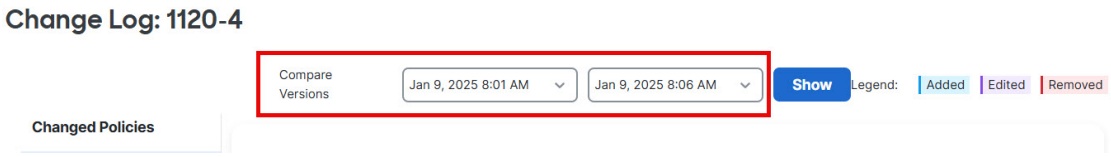

Note
The drop-down boxes also show failed deployments.

b. The **Modified By** column lists the users who have modified the policies or objects.

1. At the policy level, Firewall Management Center displays all the user names who have modified the policy.
2. At the rule level, Firewall Management Center displays the last user who has modified the rule.

c. You can also download a copy of the change log by clicking **Download Report**.

Note
- Deployment history preview is not supported for certificate enrollments, HA operations, and failed deployments.
- When a device is registered, preview is not supported for the job history record that is created.

Step 5 (Optional) Against each deployment job, click the **More** (⋮) icon and execute other actions:

- **Bookmark**—To bookmark the deployment job.
- **Edit Deployment Notes**—To edit your custom deployment notes that you added for a deployment job.
- **Generate Report**—To generate a deployment report, which can be used for auditing. This report includes job properties with preview and transcript information, and the report can be downloaded as a PDF file.

  a. Click **Generate Report** to generate a deployment report.

Figure 113: Generate Report

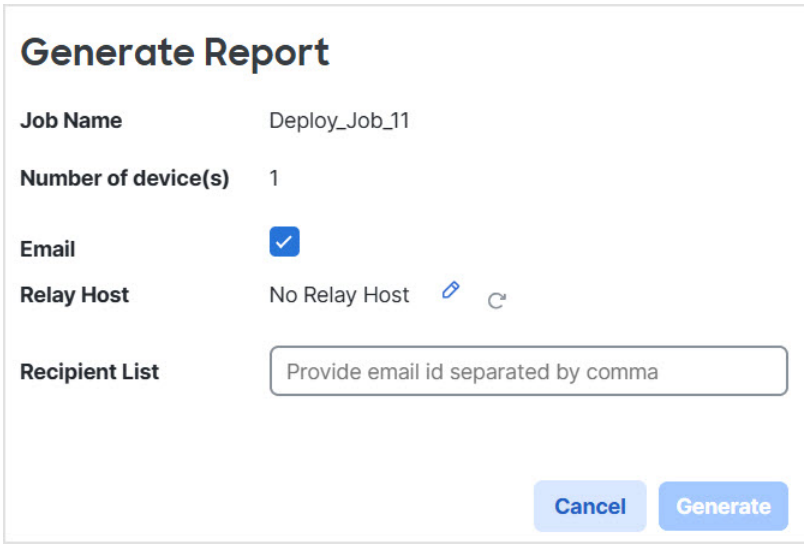

b. In the **Generate Report** popup window, check the **Email** checkbox.

c. The report can also be sent through email if mail relay host is configured. If the mail relay host in not configured, use the **Edit** (✎) icon to configure or modify the mail relay host. For more information, see *Configuring a Mail Relay Host and Notification Address* in the Cisco Secure Firewall Management Center Administration Guide.

d. In the **Recipient List**, you can enter multiple email adddresses, separated by semicolons.

e. Click **Generate** to generate the report, and this report is emailed to the recipients.

f. In the Notifications task tab, you can track the progress. After the report generation is complete, click the link in the notification task tab to download the PDF report.

Set the Number of Configuration Versions

The Firewall Management Center stores device configuration history files on the disk as configuration versions. You can specify the number of configuration versions that you want to retain for a device. This setting allows you to estimate the size of the device configuration files on the disk and keep it within the allowed limit. Reducing the number of configuration versions can reduce the backup size and improve the high-availability synchronization speed of the Firewall Management Center.

In a Firewall Management Center high-availability deployment, configuration version setting is available only on the active Firewall Management Center.

Before you begin

This feature is not supported in Version 7.4.0.

Procedure

Step 1 On the Firewall Management Center menu bar, choose **Deploy** > **Deployment History** ( ).

Step 2 Click **Deployment Setting**.

Step 3 Choose the number of configuration versions that you want to retain for a device from the **Number of Versions to Retain** drop-down list.

Note
Reducing the number of versions removes the oldest configuration versions to match the version size you have selected. You cannot roll back or preview the removed versions.

- **Maximum Permitted Disk Size**: The maximum size for storing configuration versions is 20 GB. The Firewall Management Center calculates the size of configuration versions periodically and sends a health alert if the size of the configuration versions exceeds 20 GB. To resolve the health alert, choose a **Number of Versions to Retain** for which the estimated configuration version size is less than 20 GB.
- **Current Configuration Version Size**: The size of the configuration files on the Firewall Management Center for the previous deployment.
- **Estimated Configuration Version Size**: The approximate size of the configuration files on the Firewall Management Center. It is calculated based on the number of configuration versions you chose to retain.

Step 4 Click **Save**.

Roll Back a Deployment

You can roll back a device to a previously deployed configuration. After a policy deployment, if the traffic through the device is affected in an unintended way, rollback provides an option to revert the device to the earlier state, which existed before the faulty deployment.

Rollback is a disruptive operation: all the existing connections and routes are dropped, and the traffic is disrupted.

Identifying the Disruptive Configuration

When a deployment has gone awry and caused traffic interruption in an unintended way, you should identify the change in the deployment that caused the condition and fix it so your next deployment will be successful.

See the following ways to compare configurations.

Before a Rollback

1. Choose **Deploy** > **Deployment History**, expand the last deployed job (that caused the traffic disruption), and click the **Preview** ( ).

   The preview page provides an option to compare deployments, which can be useful to identify specific changes for a deployment compared to a previous deployment.

2. After identifying the change causing the problem, rectify the configuration, and redeploy it on the device.

After a Rollback

1. After a successful rollback operation, choose **Deploy** > **Deployment**, and click the **Preview** icon next to the rolled back device.
2. View the changes between the rolled back configuration and the current changes in the Firewall Management Center that are pending deployment.
3. After identifying the change causing the problem, rectify the configuration, and redeploy it on the device.

Rollback Guidelines and Limitations

- You can roll back to any one of the last 10 versions before the currently deployed version. Rollback to versions prior to these are not supported. The rollback icon is greyed out for unsupported versions.
- You have to perform a deployment before you can roll back again.
- After you perform a rollback, the rolled back devices are marked as out-of-date on the Firewall Management Center. The changes you made to the configuration are still pending for the next deployment. To see the pending changes, choose **Deploy** > **Deployment**, and click the **Preview** icon next to the rolled back device.
- For devices with very large access lists, if the **Object Group Search** setting is disabled, then the rollback operation may take a longer duration to complete. To verify the **Object Group Search** setting, choose **Devices** > **Device Management**, and then select the device and click **Edit Advanced Settings**.
- For the Firepower 4100/9300, make sure your current Firewall Chassis Manager interface configuration is the same for any rollback versions. Otherwise, the rollback interface configuration may not match your actual interfaces.
- Rollback is not supported if the manager access interface (Manager or data interface) is different between the rollback version and the current version.
- Independent certificate enrollments are also listed as deployment jobs in the Deployment History page. However, you cannot roll back to these versions. A rollback from a deployment version created after certificate enrollments reverts the certificate associations as well. In the next deployment after a rollback, manually associate the certificates before proceeding with the deployment.
- If you upgrade the Firewall Management Center, all rollback versions from the previous software release will no longer be available for devices, even if you did not upgrade the devices.
- If you upgrade the device, you can only roll back to versions on the current software release.
- If a deployment for a device with a FlexConfig object configured with a deployment frequency set to **Once** is rolled back, then you will not be able to redeploy the object even though it is displayed as out-of-date on the Preview page. After a rollback, you will have to manually unassign and then reassign the FlexConfig object to the device before the next deployment.
- For High Availability, rollback is not supported in the following scenarios:
  - When the version you want to roll back to contains the high-availability bootstrap configuration. In other words, the deployment when you first formed high availability for the standalone devices.
  - When a device that is currently in standalone mode was part of a high availability pair in the previous deployment version.
- For clustering, see the following guidelines:

- Rollback is not supported when a device that is currently in standalone mode was part of a cluster in the previous deployment version.
- (Secure Firewall 3100/4200 and Firewall Threat Defense Virtual in a private cloud) If you change the clustering bootstrap configuration or add or delete nodes, you cannot roll back to a version prior to those changes.

Configurations Not Reverted After a Rollback

Rollback reverts all the configurations on the device except a few. See the table below for details.

Configurations that are reverted during a rollback	Configurations that are not reverted during a rollback
• All policy configurations • Interface configurations • SRU configurations • VDB configurations • LSP configurations • VPN configurations • FXOS configurations	• Snort binaries • Geo DB

Perform a Rollback

You can roll back a device to a previously deployed configuration. After a policy deployment, if the traffic through the device is affected in an unintended way, rollback provides an option to revert the device to the earlier state, which existed before the faulty deployment.

Rollback reverts the configuration only on the selected devices.

Procedure

Step 1 Choose **Deploy >Deployment History** ( ).

A list of all the previous deployment jobs is displayed in reverse chronological order.

Step 2 Click **Rollback**.

Step 3 Filter the list of devices displayed by either clicking **Job** and choosing a job from the **Selected Job** drop-down list, or by clicking **Device List**.

Step 4 (Optional) Enter the device name in the **Search Device** search box to filter the device list.

Step 5 Check the box next to the devices you want to roll back and choose the version for each device from the **Rollback Version** drop-down list.

Figure 114: Selected Job List

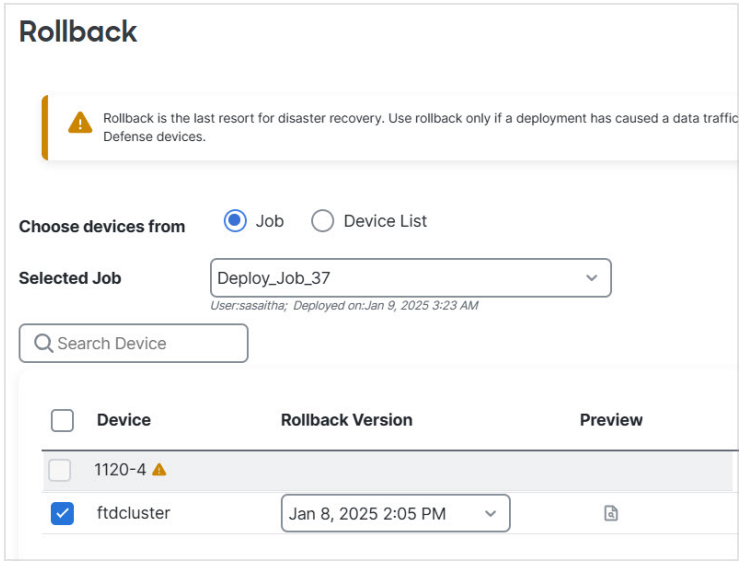

Figure 115: Device List

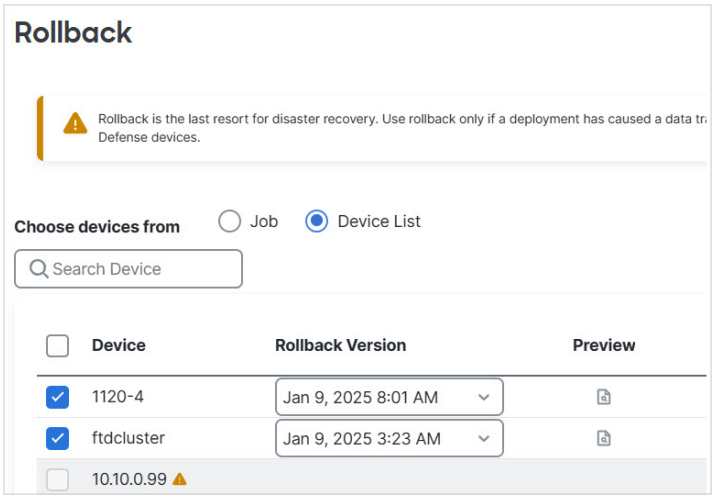

The job names and associated deployment notes are also listed for a particular rollback version.

Step 6 (Optional) Click the **Preview** ( ) to view the changes deployed in the selected version.

Step 7 Click **Rollback**.

What to do next

To know the status of the rollback, choose **Deploy** > **Deployment**. You can view the rollback status next to the device name.

View the Deployment Rollback Transcript

Rollback transcript is a written version of the commands that are sent to the device, along with the responses returned from the device. If a rollback operation has failed, the transcript in the **Deploy > Deployment History** page provides the reason for the failure. However, to know the CLI commands executed for a successful rollback operation, follow the steps given below after a rollback operation has completed. Note that this information is available only until the next deployment.

Note The CLI command information is available after a rollback is complete and is available only until the next deployment. The first deployment after a rollback operation erases all the rollback-related information.

Note For any rollback deployment, the Deployment Notes are updated automatically as Rollback Job. In the **Deployment History** page, the user can filter the Rollback Jobs easily using the Search option.

Procedure

Step 1 On the Secure Firewall Management Center menu bar, choose **System** (⚙) > **Health** > **Monitor**.

Step 2 Select the device, which was rolled back, from the left pane.

Step 3 Click **View System & Troubleshooting Details** link.

Step 4 Click **Advanced Troubleshooting**.

Step 5 Click **Threat Defense CLI**.

Step 6 Select **show** from the **Command** drop-down box.

Step 7 Enter `running` in the **Parameter** field.

Step 8 Click **Execute**.

Download Policy Changes Report for Multiple Devices

Download reports on the policy and object changes made since your last deployment for multiple Firewall Threat Defense devices. You can download the reports in the form of a zip file that contains the following reports:

- A pending changes report for each device, that previews the additions, updates, or deletions in the policy, or the objects that are to be deployed on the device. For more information, see Deploy Configuration Changes, on page 259 and Deployment Preview.
- A consolidated report that categorizes each device based on the report status.

Procedure

Step 1 Choose **Deploy** > **Advanced Deploy**.

Step 2 Check the check box next to the devices for which you want to generate a pending policy changes report, and then click **Pending Changes Reports**.

Step 3 Click **Pending Changes Reports**. Reports are generated in the background.

Step 4 On the Firewall Management Center menu bar, choose **Notifications** > **Tasks** to view the report generation task.

When the report request task is complete, the download link appears within the task notification.

Step 5 Click the **Download Report(s)** link to download the reports.

Compare Policies

To review policy changes for compliance with your organization's standards or to optimize system performance, you can examine the differences between two policies or between a saved policy and the running configuration.

You can compare the following policy types:

- DNS
- File
- Health
- Identity
- Network Analysis
- SSL

The comparison view displays both policies in a side-by-side format. Differences between the two policies are highlighted:

- Blue indicates that the highlighted setting is different in the two policies, and the difference is noted in red text.
- Green indicates that the highlighted setting appears in one policy but not the other.

Before you begin

You can compare policies only if you have access rights and any required licenses for the specific policy, and you are in the correct domain for configuring the policy.

Procedure

Step 1 Access the management page for the policy you want to compare:

- DNS—**Policies** > **Access Control heading** > **DNS**
- File—**Policies** > **Access Control heading** > **Malware & File**
- Health—**System** (⚙) > **Health** > **Policy**
- Identity—**Policies** > **Access Control heading** > **Identity**
- Network Analysis—**Policies** > **Access Control heading** > **Access Control**, and then click **Network Analysis Policy** or **Policies** > **Access Control heading** > **Intrusion**, and then click **Network Analysis Policies**.

  Note
  If your custom user role limits access to the first path listed here, use the second path to access the policy.

- SSL—**Policies** > **Access Control heading** > **Decryption**

Step 2 Click **Compare Policies**.

Step 3 From the **Compare Against** drop-down list, choose the type of comparison you want to make:

- To compare two different policies, choose **Other Policy**.
- To compare two revisions of the same policy, choose **Other Revision**.
- To compare another policy to the currently active policy, choose **Running Configuration**.

Step 4 Depending on the comparison type you choose, you have the following choices:

- If you are comparing two different policies, choose the policies you want to compare from the **Policy A** and **Policy B** drop-down lists.
- If you are comparing the running configuration to another policy, choose the second policy from the **Policy B** drop-down list.

Step 5 Click **OK**.

Step 6 Review the comparison results:

- Comparison Viewer—To use the comparison viewer to navigate individually through policy differences, click **Previous** or **Next** above the title bar.
- Comparison Report—To generate a PDF report that lists the differences between the two policies, click **Comparison Report**.

Generate Current Policy Reports

For most policies, you can generate two kinds of reports. A report on a single policy provides details on the policy's current saved configuration, while a comparison report lists only the differences between two policies. You can generate a single-policy report for all policy types except health.

Note Intrusion policy reports combine the settings in the base policy with the settings of the policy layers, and make no distinction between which settings originated in the base policy or policy layer.

Before you begin

You can generate policy reports only if you have access rights and any required licenses for the specific policy, and you are in the correct domain for configuring the policy.

Procedure

Step 1 Access the management page for the policy for which you want to generate a report:

- Access Control—**Policies** > **Access Control heading** > **Access Control**
- DNS—**Policies** > **Access Control heading** > **DNS**
- File—**Policies** > **Access Control heading** > **Malware & File**
- Health—**System** (⚙) > **Health** > **Policy**
- Identity—**Policies** > **Access Control heading** > **Identity**
- Intrusion—**Policies** > **Access Control heading** > **Intrusion**
- NAT—**Devices** > **NAT**
- Network Analysis—**Policies** > **Access Control heading** > **Access Control**, and then click **Network Analysis Policy** or **Policies** > **Access Control heading** > **Intrusion**, and then click **Network Analysis Policies**

  Note
  If your custom user role limits access to the first path listed here, use the second path to access the policy.

- SSL—**Policies** > **Access Control heading** > **Decryption**

Step 2 Click **Report** (📋) next to the policy for which you want to generate a report.

History for Configuration Deployment

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Set the number of deployment history files to retain for device rollback.	7.2.6 7.4.1	Any	You can now set the number of deployment history files to retain for device rollback, up to ten (the default). This can help you save disk space on the Firewall Management Center. New/modified screens: **Deploy** > **Deployment History** > **Deployment Setting** > **Configuration Version Setting** Version restrictions: Not supported with Firewall Management Center Version 7.3.x or 7.4.0.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
View and generate reports on configuration changes since your last deployment.	7.2.6 7.4.1	Any	You can generate, view, and download (as a zip file) the following reports on configuration changes since your last deployment: • A policy changes report for each device that previews the additions, changes, or deletions in the policy, or the objects that are to be deployed on the device. • A consolidated report that categorizes each device based on the status of policy changes report generation. This is especially useful after you upgrade either the Firewall Management Center or Firewall Threat Defense, so that you can see the changes made by the upgrade before you deploy. New/modified screens: **Deploy** > **Advanced Deploy**. Version restrictions: Not supported with Firewall Management Center Version 7.3.x or 7.4.0.
Generate and email a report when you deploy configuration changes.	7.2	Any	You can now generate a report for any deployment. New/modified screens: **Deploy** > **Deployment History icon** > **More** > **Generate Report**
Deployment preview and user information in preview.	7.0	Any	The Deployment page has the following newly added features: • On the **Deployment** page, the **Modified By** column lists the users who have modified the policies against each policy listing. • Filter support for deployment – The Filter icon provided on the **Deployment** page provides an option to filter the device listings that are pending deployment. The Filter icon provides options to filter the listings based on selected devices and user names. • Deployment History Preview – Click **Preview** to view the policy and object changes deployed on the device versus the previously deployed version. In the deployment history, the last 10 successful deployments, the last five failed deployments, and last five rollback deployments are captured. • Deployment Notes – The **Deployment Notes** are custom and optional notes that a user can add as part of deployment. You can view the **Deployment Notes** column in the **Deployment History** page. • Deployment rollback is available for Snort 3 policies as well.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Roll back deployment on Firewall Threat Defense devices.	6.7	6.7	Rollback is a deployment functionality provided to remove the existing deployment on Firewall Threat Defense devices and to reconfigure the device with the previously deployed configuration. New/modified pages: The **Deploy** > **Deployment History** page provides a new Rollback column with the rollback icons. Similar rollback icons can also be found when the jobs are expanded, to initiate rollback at the device level.
New deployment web interface.	6.6	Any	The **Deploy** button on the Firewall Management Center menu bar is changed to **Deploy** menu. There are two new sub-menu options under it. These are **Deployment** and **Deployment History**. The Deployment page has undergone an improvement along with newly added features, and the new Deployment History page provides a legend of all the previous deployments. The Deployment page has the following newly added features: • Deployment status - On the Deployment page, the Status column provides the deployment status for each device. • Deployment estimate - The **Estimate** link is available on the Deployment page after you select a device, a policy, or a configuration. The **Estimate** link provides an estimate of the deployment duration once clicked. • Deployment preview - Preview provides a snapshot of all the policy and object changes to be deployed on the device. The policy changes include the new policies, changes in the existing policies, and the deleted policies. The object changes include the added and modified objects which are used in policies. • Selective policy deployment - Firewall Management Center allows you to select a specific policy within the list of all the changes on the device that are due for deployment and deploy only the selected policy.

PART II

Device Operations

- Transparent or Routed Firewall Mode, on page 285
- Logical Devices on the Firepower 4100/9300, on page 297
- Multi-Instance Mode for the Secure Firewall 3100/4200, on page 355
- High Availability, on page 423
- Clustering for the Secure Firewall 3100/4200, on page 457
- Clustering for Threat Defense Virtual in a Private Cloud, on page 525
- Clustering for Threat Defense Virtual in a Public Cloud, on page 579
- Clustering for the Firepower 4100/9300, on page 699

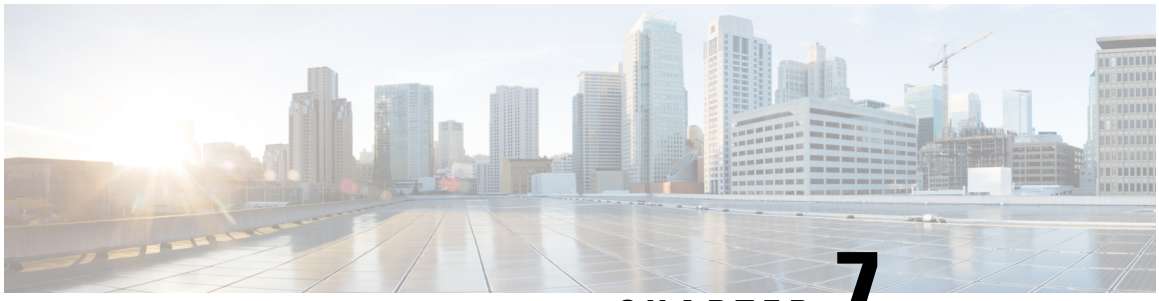

CHAPTER 7

Transparent or Routed Firewall Mode

This chapter describes how to set the firewall mode to routed or transparent, as well as how the firewall works in each firewall mode.

Note The firewall mode only affects regular firewall interfaces, and not IPS-only interfaces such as inline sets or passive interfaces. IPS-only interfaces can be used in both firewall modes. See Inline Sets and Passive Interfaces, on page 893 for more information about IPS-only interfaces. Inline sets might be familiar to you as "transparent inline sets," but the inline interface type is unrelated to the transparent firewall mode described in this chapter or the firewall-type interfaces.

Caution

- Use FTD CLI commands to set "Firewall Mode."

- About the Firewall Mode, on page 285
- Default Settings, on page 293
- Guidelines for Firewall Mode, on page 293
- Set the Firewall Mode, on page 294

About the Firewall Mode

The Firewall Threat Defense supports two firewall modes for regular firewall interfaces: Routed Firewall mode and Transparent Firewall mode.

About Routed Firewall Mode

In routed mode, the Firewall Threat Defense device is considered to be a router hop in the network. Each interface that you want to route between is on a different subnet.

With Integrated Routing and Bridging, you can use a "bridge group" where you group together multiple interfaces on a network, and the Firewall Threat Defense device uses bridging techniques to pass traffic between the interfaces. Each bridge group includes a Bridge Virtual Interface (BVI) to which you assign an IP address on the network. The Firewall Threat Defense device routes between BVIs and regular routed interfaces. If you do not need clustering or EtherChannel member interfaces, you might consider using routed

mode instead of transparent mode. In routed mode, you can have one or more isolated bridge groups like in transparent mode, but also have normal routed interfaces as well for a mixed deployment.

About Transparent Firewall Mode

Traditionally, a firewall is a routed hop and acts as a default gateway for hosts that connect to one of its screened subnets. A transparent firewall, on the other hand, is a Layer 2 firewall that acts like a “bump in the wire,” or a “stealth firewall,” and is not seen as a router hop to connected devices. However, like any other firewall, access control between interfaces is controlled, and all of the usual firewall checks are in place.

Layer 2 connectivity is achieved by using a "bridge group" where you group together the inside and outside interfaces for a network, and the Firewall Threat Defense device uses bridging techniques to pass traffic between the interfaces. Each bridge group includes a Bridge Virtual Interface (BVI) to which you assign an IP address on the network. You can have multiple bridge groups for multiple networks. In transparent mode, these bridge groups cannot communicate with each other.

Using the Transparent Firewall in Your Network

The Firewall Threat Defense device connects the same network between its interfaces. Because the firewall is not a routed hop, you can easily introduce a transparent firewall into an existing network.

The following figure shows a typical transparent firewall network where the outside devices are on the same subnet as the inside devices. The inside router and hosts appear to be directly connected to the outside router.

Figure 116: Transparent Firewall Network

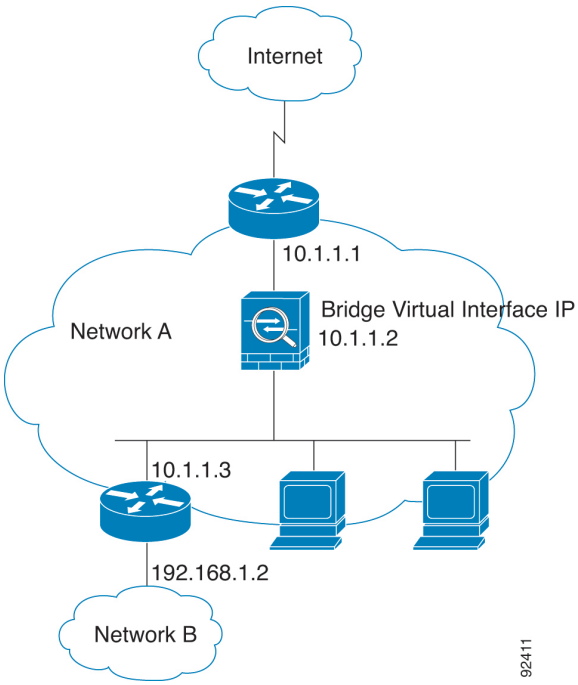

Passing Traffic For Routed-Mode Features

For features that are not directly supported on the transparent firewall, you can allow traffic to pass through so that upstream and downstream routers can support the functionality. For example, by using an access rule,

you can allow DHCP traffic (instead of the unsupported DHCP relay feature) or multicast traffic such as that created by IP/TV. You can also establish routing protocol adjacencies through a transparent firewall; you can allow OSPF, RIP, EIGRP, or BGP traffic through based on an access rule. Likewise, protocols like HSRP or VRRP can pass through the Firewall Threat Defense device.

About Bridge Groups

A bridge group is a group of interfaces that the Firewall Threat Defense device bridges instead of routes. Bridge groups are supported in both transparent and routed firewall mode. Like any other firewall interfaces, access control between interfaces is controlled, and all of the usual firewall checks are in place.

Bridge Virtual Interface (BVI)

Each bridge group includes a Bridge Virtual Interface (BVI). The Firewall Threat Defense device uses the BVI IP address as the source address for packets originating from the bridge group. The BVI IP address must be on the same subnet as the bridge group member interfaces. The BVI does not support traffic on secondary networks; only traffic on the same network as the BVI IP address is supported.

In transparent mode: Only bridge group member interfaces are named and can be used with interface-based features.

In routed mode: The BVI acts as the gateway between the bridge group and other routed interfaces. To route between bridge groups/routed interfaces, you must name the BVI. For some interface-based features, you can use the BVI itself:

- DHCPv4 server—Only the BVI supports the DHCPv4 server configuration.
- Static routes—You can configure static routes for the BVI; you cannot configure static routes for the member interfaces.
- Syslog server and other traffic sourced from the Firewall Threat Defense device—When specifying a syslog server (or SNMP server, or other service where the traffic is sourced from the Firewall Threat Defense device), you can specify either the BVI or a member interface.

If you do not name the BVI in routed mode, then the Firewall Threat Defense device does not route bridge group traffic. This configuration replicates transparent firewall mode for the bridge group. If you do not need clustering or EtherChannel member interfaces, you might consider using routed mode instead. In routed mode, you can have one or more isolated bridge groups like in transparent mode, but also have normal routed interfaces as well for a mixed deployment.

Bridge Groups in Transparent Firewall Mode

Bridge group traffic is isolated from other bridge groups; traffic is not routed to another bridge group within the Firewall Threat Defense device, and traffic must exit the Firewall Threat Defense device before it is routed by an external router back to another bridge group in the Firewall Threat Defense device. Although the bridging functions are separate for each bridge group, many other functions are shared between all bridge groups. For example, all bridge groups share a syslog server or AAA server configuration.

You can include multiple interfaces per bridge group. See Guidelines for Firewall Mode, on page 293 for the exact number of bridge groups and interfaces supported. If you use more than 2 interfaces per bridge group, you can control communication between multiple segments on the same network, and not just between inside and outside. For example, if you have three inside segments that you do not want to communicate with each other, you can put each segment on a separate interface, and only allow them to communicate with the outside interface. Or you can customize the access rules between interfaces to allow only as much access as desired.

The following figure shows two networks connected to the Firewall Threat Defense device, which has two bridge groups.

Figure 117: Transparent Firewall Network with Two Bridge Groups

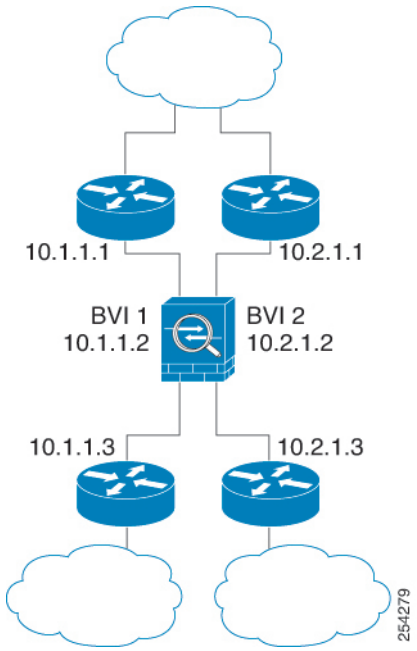

Bridge Groups in Routed Firewall Mode

Bridge group traffic can be routed to other bridge groups or routed interfaces. You can choose to isolate bridge group traffic by not assigning a name to the BVI interface for the bridge group. If you name the BVI, then the BVI participates in routing like any other regular interface.

One use for a bridge group in routed mode is to use extra interfaces on the Firewall Threat Defense instead of an external switch. For example, the default configuration for some devices include an outside interface as a regular interface, and then all other interfaces assigned to the inside bridge group. Because the purpose of this bridge group is to replace an external switch, you need to configure an access policy so all bridge group interfaces can freely communicate.

Figure 118: Routed Firewall Network with an Inside Bridge Group and an Outside Routed Interface

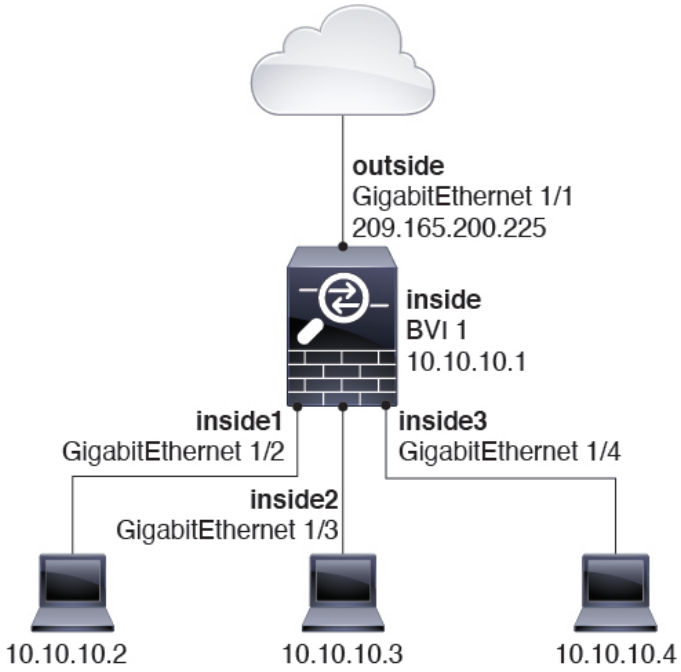

Allowing Layer 3 Traffic

- Unicast IPv4 and IPv6 traffic requires an access rule to be allowed through the bridge group.
- ARPs are allowed through the bridge group in both directions without an access rule. ARP traffic can be controlled by ARP inspection.
- IPv6 neighbor discovery and router solicitation packets can be passed using access rules.
- Broadcast and multicast traffic can be passed using access rules.

Allowed MAC Addresses

The following destination MAC addresses are allowed through the bridge group if allowed by your access policy (see Allowing Layer 3 Traffic, on page 289). Any MAC address not on this list is dropped.

- TRUE broadcast destination MAC address equal to FFFF.FFFF.FFFF
- IPv4 multicast MAC addresses from 0100.5E00.0000 to 0100.5EFE.FFFF
- IPv6 multicast MAC addresses from 3333.0000.0000 to 3333.FFFF.FFFF
- BPDU multicast address equal to 0100.0CCC.CCCD

BPDU Handling

To prevent loops using the Spanning Tree Protocol, BPDUs are passed by default.

By default BPDUs are also forwarded for advanced inspection, which is unnecessary for this type of packet, and which can cause problems if they are blocked due to an inspection restart, for example. We recommend

that you always exempt BPDUs from advanced inspection. To do so, use FlexConfig to configure an EtherType ACL that trusts BPDUs and exempts them from advanced inspection on each member interface. See FlexConfig Policies, on page 2667.

The FlexConfig object should deploy the following commands, where you replace <if-name> with an interface name. Add as many access-group commands as needed to cover each bridge group member interface on the device. You can also choose a different name for the ACL.

```
access-list permit-bpdu ethertype trust bpdu
access-group permit-bpdu in interface <if-name>
```

MAC Address vs. Route Lookups

For traffic within a bridge group, the outgoing interface of a packet is determined by performing a destination MAC address lookup instead of a route lookup.

Route lookups, however, are necessary for the following situations:

- Traffic originating on the Firewall Threat Defense device—Add a default/static route on the Firewall Threat Defense device for traffic destined for a remote network where a syslog server, for example, is located.
- Voice over IP (VoIP) and TFTP traffic, and the endpoint is at least one hop away—Add a static route on the Firewall Threat Defense device for traffic destined for the remote endpoint so that secondary connections are successful. The Firewall Threat Defense device creates a temporary "pinhole" in the access control policy to allow the secondary connection; and because the connection might use a different set of IP addresses than the primary connection, the Firewall Threat Defense device needs to perform a route lookup to install the pinhole on the correct interface.

  Affected applications include:

  - H.323
  - RTSP
  - SIP
  - Skinny (SCCP)
  - SQL*Net
  - SunRPC
  - TFTP

- Traffic at least one hop away for which the Firewall Threat Defense device performs NAT—Configure a static route on the Firewall Threat Defense device for traffic destined for the remote network. You also need a static route on the upstream router for traffic destined for the mapped addresses to be sent to the Firewall Threat Defense device.

  This routing requirement is also true for embedded IP addresses for VoIP and DNS with NAT enabled, and the embedded IP addresses are at least one hop away. The Firewall Threat Defense device needs to identify the correct egress interface so it can perform the translation.

Figure 119: NAT Example: NAT within a Bridge Group

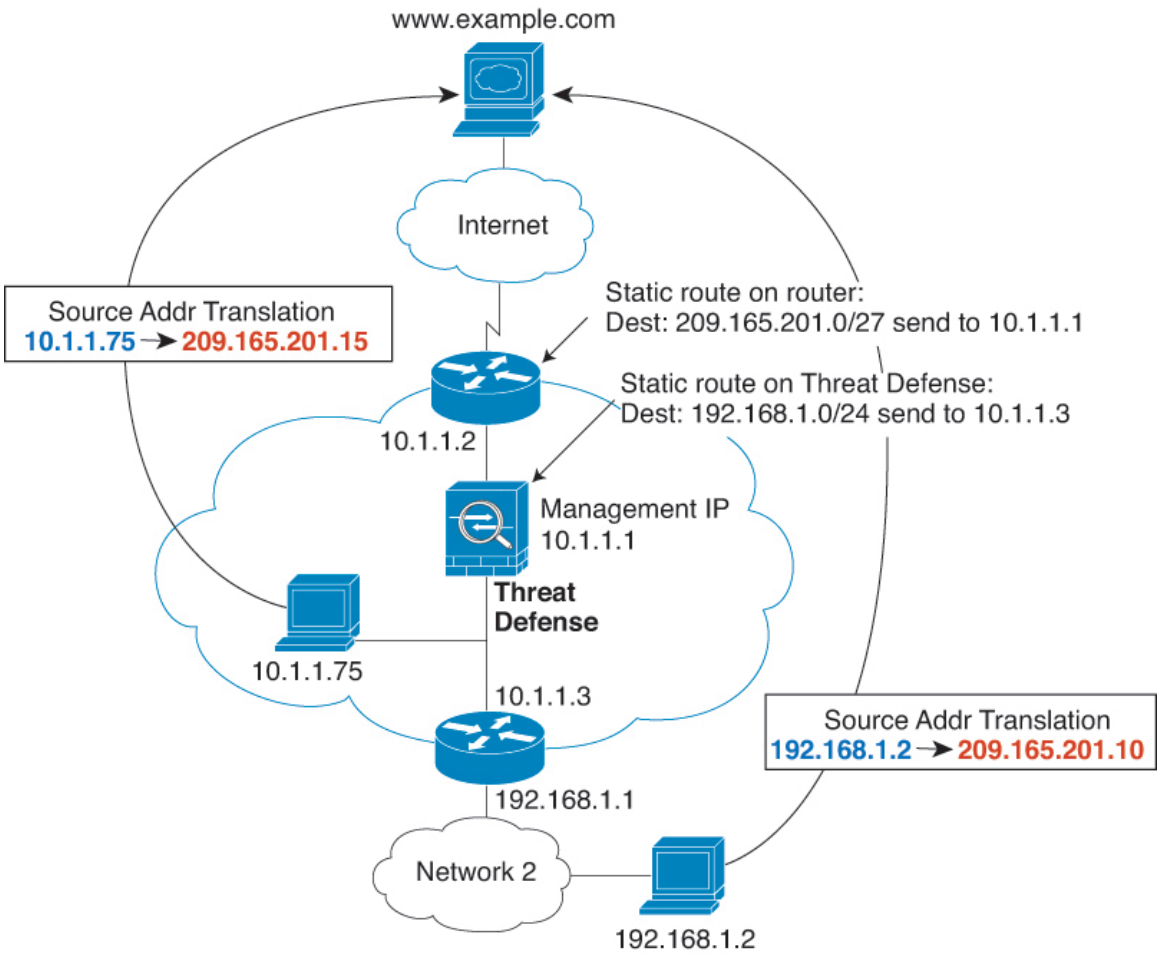

Unsupported Features for Bridge Groups in Transparent Mode

The following table lists the features are not supported in bridge groups in transparent mode.

Table 15: Unsupported Features in Transparent Mode

Feature	Description
Dynamic DNS	—
DHCPv6 stateless server	Only the DHCPv4 server is supported on bridge group member interfaces.
DHCP relay	The transparent firewall can act as a DHCPv4 server, but it does not support DHCP relay. DHCP relay is not required because you can allow DHCP traffic to pass through using two access rules: one that allows DCHP requests from the inside interface to the outside, and one that allows the replies from the server in the other direction.

Feature	Description
Dynamic routing protocols	You can, however, add static routes for traffic originating on the Firewall Threat Defense device for bridge group member interfaces. You can also allow dynamic routing protocols through the Firewall Threat Defense device using an access rule.
Multicast IP routing	You can allow multicast traffic through the Firewall Threat Defense device by allowing it in an access rule.
QoS	—
VPN termination for through traffic	The transparent firewall supports site-to-site VPN tunnels for management connections only on bridge group member interfaces. It does not terminate VPN connections for traffic through the Firewall Threat Defense device. You can pass VPN traffic through the Firewall Threat Defense device using an access rule, but it does not terminate non-management connections.

Unsupported Features for Bridge Groups in Routed Mode

The following table lists the features are not supported in bridge groups in routed mode.

Table 16: Unsupported Features in Routed Mode

Feature	Description
EtherChannel member interfaces	Only physical interfaces, redundant interfaces, and subinterfaces are supported as bridge group member interfaces. Management interfaces are also not supported.
Clustering	Bridge groups are not supported in clustering.
Dynamic DNS	—
DHCPv6 stateless server	Only the DHCPv4 server is supported on BVIs.
DHCP relay	The routed firewall can act as a DHCPv4 server, but it does not support DHCP relay on BVIs or bridge group member interfaces.
Dynamic routing protocols	You can, however, add static routes for BVIs. You can also allow dynamic routing protocols through the Firewall Threat Defense device using an access rule. Non-bridge group interfaces support dynamic routing.
Multicast IP routing	You can allow multicast traffic through the Firewall Threat Defense device by allowing it in an access rule. Non-bridge group interfaces support multicast routing.
QoS	Non-bridge group interfaces support QoS.

Feature	Description
VPN termination for through traffic	You cannot terminate a VPN connection on the BVI. Non-bridge group interfaces support VPN. Bridge group member interfaces support site-to-site VPN tunnels for management connections only. It does not terminate VPN connections for traffic through the Firewall Threat Defense device. You can pass VPN traffic through the bridge group using an access rule, but it does not terminate non-management connections.

Default Settings

Bridge Group Defaults

By default, all ARP packets are passed within the bridge group.

Guidelines for Firewall Mode

Bridge Group Guidelines (Transparent and Routed Mode)

- You can create up to 250 bridge groups, with 64 interfaces per bridge group.
- Each directly-connected network must be on the same subnet.
- The Firewall Threat Defense device does not support traffic on secondary networks; only traffic on the same network as the BVI IP address is supported.
- An IP address for the BVI is required for each bridge group for to-the-device and from-the-device management traffic, as well as for data traffic to pass through the Firewall Threat Defense device. For IPv4 traffic, specify an IPv4 address. For IPv6 traffic, specify an IPv6 address.
- You can only configure IPv6 addresses manually.
- The BVI IP address must be on the same subnet as the connected network. You cannot set the subnet to a host subnet (255.255.255.255).
- Management interfaces are not supported as bridge group members.
- For multi-instance mode, shared interfaces are not supported for bridge group member interfaces (in transparent mode or routed mode).
- For the Firewall Threat Defense Virtual on VMware with bridged ixgbevf interfaces, transparent mode is not supported, and bridge groups are not supported in routed mode.
- For the Firepower 1010 and Secure Firewall 1210/20, you cannot mix logical VLAN interfaces and physical firewall interfaces in the same bridge group.
- For the Firepower 4100/9300, data-sharing interfaces are not supported as bridge group members.
- In transparent mode, you must use at least 1 bridge group; data interfaces must belong to a bridge group.

- In transparent mode, do not specify the BVI IP address as the default gateway for connected devices; devices need to specify the router on the other side of the Firewall Threat Defense as the default gateway.
- In transparent mode, the *default* route, which is required to provide a return path for management traffic, is only applied to management traffic from one bridge group network. This is because the default route specifies an interface in the bridge group as well as the router IP address on the bridge group network, and you can only define one default route. If you have management traffic from more than one bridge group network, you need to specify a regular static route that identifies the network from which you expect management traffic.
- Transparent mode is not supported on threat defense virtual instances deployed on Amazon Web Services, Microsoft Azure, Google Cloud Platform, and Oracle Cloud Infrastructure.
- In routed mode, to route between bridge groups and other routed interfaces, you must name the BVI.
- In routed mode, Firewall Threat Defense-defined EtherChannel interfaces are not supported as bridge group members. EtherChannels on the Firepower 4100/9300 can be bridge group members.
- Bidirectional Forwarding Detection (BFD) echo packets are not allowed through the Firewall Threat Defense when using bridge group members. If there are two neighbors on either side of the Firewall Threat Defense running BFD, then the Firewall Threat Defense will drop BFD echo packets because they have the same source and destination IP address and appear to be part of a LAND attack.

Set the Firewall Mode

You can set the firewall mode when you perform the initial system setup at the CLI. We recommend setting the firewall mode during setup because changing the firewall mode erases your configuration to ensure you do not have incompatible settings. If you need to change the firewall mode later, you must do so from the CLI.

Procedure

Step 1 Unregister the Firewall Threat Defense device from the Firewall Management Center.

You cannot change the mode until you unregister the device.

a) Performance profiles are retained across firewall mode changes and can't be changed after you re-register. You must set the performance profile to Default. See the **Performance Profile** page in your platform settings policy (choose **Devices** > **Platform Settings**, and then click **Edit** (✎) for your policy).

In addition, if you enabled Universal Zero Trust Network Access (universal ZTNA) for the device, which also enabled a special universal ZTNA performance profile, you need to disable it in Security Cloud Control. See **Device Management** > **Universal zero trust access settings**.

Deploy the configuration to make these changes.

b) Choose **Devices** > **Device Management**.

c) Next to the device you want to unregister, click **More** (⋮), and then click **Delete**.

Step 2 Access the Firewall Threat Defense device CLI, preferably from the console port.

Step 3 Change the firewall mode:

configure firewall [**routed** | **transparent**]

Example:

```
> configure firewall transparent
This will destroy the current interface configurations, are you sure that you want to
proceed? [y/N] y
The firewall mode was changed successfully.
```

Step 4 Re-register with the Firewall Management Center. See Complete the Firewall Threat Defense Initial Configuration Using the CLI, on page 19 and Add a Device Using a Registration Key—Basic Configuration, on page 31.

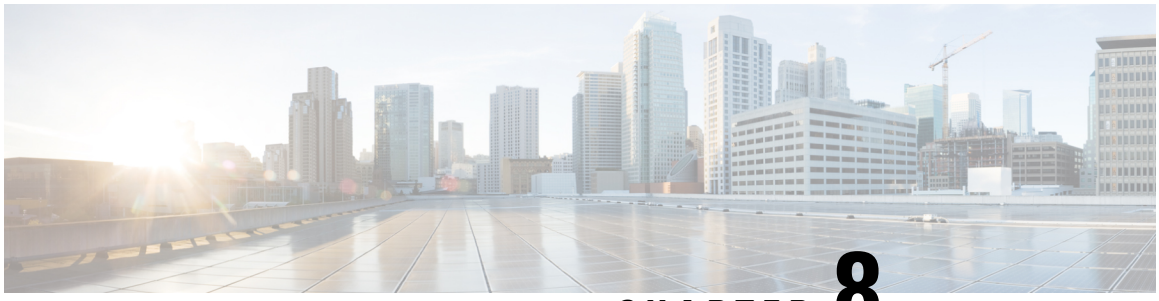

CHAPTER 8

Logical Devices on the Firepower 4100/9300

The Firepower 4100/9300 is a flexible security platform on which you can install one or more *logical devices*. Before you can add the Firewall Threat Defense to the Firewall Management Center, you must configure chassis interfaces, add a logical device, and assign interfaces to the device on the Firepower 4100/9300 chassis using the Secure Firewall Chassis Manager or the FXOS CLI. This chapter describes basic interface configuration and how to add a standalone or High Availability logical device using the Secure Firewall Chassis Manager. To add a clustered logical device, see Clustering for the Firepower 4100/9300, on page 699. To use the FXOS CLI, see the FXOS CLI configuration guide. For more advanced FXOS procedures and troubleshooting, see the FXOS configuration guide.

- About Interfaces, on page 297
- About Logical Devices, on page 311
- Licenses for Container Instances, on page 320
- Requirements and Prerequisites for Logical Devices, on page 321
- Guidelines and Limitations for Logical Devices, on page 328
- Configure Interfaces, on page 332
- Configure Logical Devices, on page 336
- History for Logical Devices, on page 348

About Interfaces

The Firepower 4100/9300 chassis supports physical interfaces, VLAN subinterfaces for container instances, and EtherChannel (port-channel) interfaces. EtherChannel interfaces can include up to 16 member interfaces of the same type.

Chassis Management Interface

The chassis management interface is used for management of the FXOS Chassis by SSH or Firewall Chassis Manager. This interface appears at the top of the **Interfaces** tab as **MGMT**, and you can only enable or disable this interface on the **Interfaces** tab. This interface is separate from the mgmt-type interface that you assign to the logical devices for application management.

To configure parameters for this interface, you must configure them from the CLI. To view information about this interface in the FXOS CLI, connect to local management and show the management port:

Firepower # **connect local-mgmt**

Firepower(local-mgmt) # **show mgmt-port**

Note that the chassis management interface remains up even if the physical cable or SFP module are unplugged, or if the **mgmt-port shut** command is performed, or if the logical device is offline.

Note The chassis management interface does not support jumbo frames.

Interface Types

Physical interfaces, VLAN subinterfaces for container instances, and EtherChannel (port-channel) interfaces can be one of the following types:

- Data—Use for regular data. Data interfaces cannot be shared between logical devices, and logical devices cannot communicate over the backplane to other logical devices. For traffic on Data interfaces, all traffic must exit the chassis on one interface and return on another interface to reach another logical device.
- Data-sharing—Use for regular data. Only supported with container instances, these data interfaces can be shared by one or more logical devices/container instances (Firewall Threat Defense-using-Firewall Management Center only). Each container instance can communicate over the backplane with all other instances that share this interface. Shared interfaces can affect the number of container instances you can deploy. Shared interfaces are not supported for bridge group member interfaces (in transparent mode or routed mode), inline sets, passive interfaces, clusters, or failover links.
- Mgmt—Use to manage application instances. These interfaces can be shared by one or more logical devices to access external hosts; logical devices cannot communicate over this interface with other logical devices that share the interface. You can only assign one management interface per logical device. Depending on your application and manager, you can later enable management from a data interface; but you must assign a Management interface to the logical device even if you don't intend to use it after you enable data management. For information about the separate chassis management interface, see Chassis Management Interface, on page 297.

Note Mgmt interface change will cause reboot of the logical device, for example one change mgmt from e1/1 to e1/2 will cause the logical device to reboot to apply the new management.

- Eventing—Use as a secondary management interface for Firewall Threat Defense-using-Firewall Management Center devices. To use this interface, you must configure its IP address and other parameters at the Firewall Threat Defense CLI. For example, you can separate management traffic from events (such as web events). See the management center configuration guide for more information. Eventing interfaces can be shared by one or more logical devices to access external hosts; logical devices cannot communicate over this interface with other logical devices that share the interface. If you later configure a data interface for management, you cannot use a separate eventing interface.

Note A virtual Ethernet interface is allocated when each application instance is installed. If the application does not use an eventing interface, then the virtual interface will be in an admin down state.

```
Firepower # show interface Vethernet775
Firepower # Vethernet775 is down (Administratively down)
Bound Interface is Ethernet1/10
Port description is server 1/1, VNIC ext-mgmt-nic5
```

- Cluster—Use as the cluster control link for a clustered logical device. By default, the cluster control link is automatically created on Port-channel 48. The Cluster type is only supported on EtherChannel interfaces. For multi-instance clustering, you cannot share a Cluster-type interface across devices. You can add VLAN subinterfaces to the Cluster EtherChannel to provide separate cluster control links per cluster. If you add subinterfaces to a Cluster interface, you cannot use that interface for a native cluster. The Firewall Device Manager and Security Cloud Control does not support clustering.

Note This chapter discusses *FXOS* VLAN subinterfaces only. You can separately create subinterfaces within the Firewall Threat Defense application. See FXOS Interfaces vs. Application Interfaces, on page 300 for more information.

See the following table for interface type support for the Firewall Threat Defense and ASA applications in standalone and cluster deployments.

Table 17: Interface Type Support

Application		Data	Data: Subinterface	Data-Sharing	Data-Sharing: Subinterface	Mgmt	Eventing	Cluster (EtherChannel only)	Cluster: Subinterface
Firewall Threat Defense	Standalone Native Instance	**Yes**	—	—	—	**Yes**	**Yes**	—	—
	Standalone Container Instance	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	—	—
	Cluster Native Instance	**Yes** (EtherChannel only for inter-chassis cluster)	—	—	—	**Yes**	**Yes**	**Yes**	—
	Cluster Container Instance	**Yes** (EtherChannel only for inter-chassis cluster)	—	—	—	**Yes**	**Yes**	**Yes**	**Yes**

Application		Data	Data: Subinterface	Data-Sharing	Data-Sharing: Subinterface	Mgmt	Eventing	Cluster (EtherChannel only)	Cluster: Subinterface
ASA	Standalone Native Instance	**Yes**	—	—	—	**Yes**	—	**Yes**	—
	Cluster Native Instance	**Yes** (EtherChannel only for inter-chassis cluster)	—	—	—	**Yes**	—	**Yes**	—

FXOS Interfaces vs. Application Interfaces

The Firepower 4100/9300 manages the basic Ethernet settings of physical interfaces, VLAN subinterfaces for container instances, and EtherChannel (port-channel) interfaces. Within the application, you configure higher level settings. For example, you can only create EtherChannels in FXOS; but you can assign an IP address to the EtherChannel within the application.

The following sections describe the interaction between FXOS and the application for interfaces.

VLAN Subinterfaces

For all logical devices, you can create VLAN subinterfaces within the application.

For container instances in standalone mode only, you can *also* create VLAN subinterfaces in FXOS. Multi-instance clusters do not support subinterfaces in FXOS except on the Cluster-type interface. Application-defined subinterfaces are not subject to the FXOS limit. Choosing in which operating system to create subinterfaces depends on your network deployment and personal preference. For example, to share a subinterface, you must create the subinterface in FXOS. Another scenario that favors FXOS subinterfaces comprises allocating separate subinterface groups on a single interface to multiple instances. For example, you want to use Port-channel1 with VLAN 2–11 on instance A, VLAN 12–21 on instance B, and VLAN 22–31 on instance C. If you create these subinterfaces within the application, then you would have to share the parent interface in FXOS, which may not be desirable. See the following illustration that shows the three ways you can accomplish this scenario:

Figure 120: VLANs in FXOS vs. the Application for Container Instances

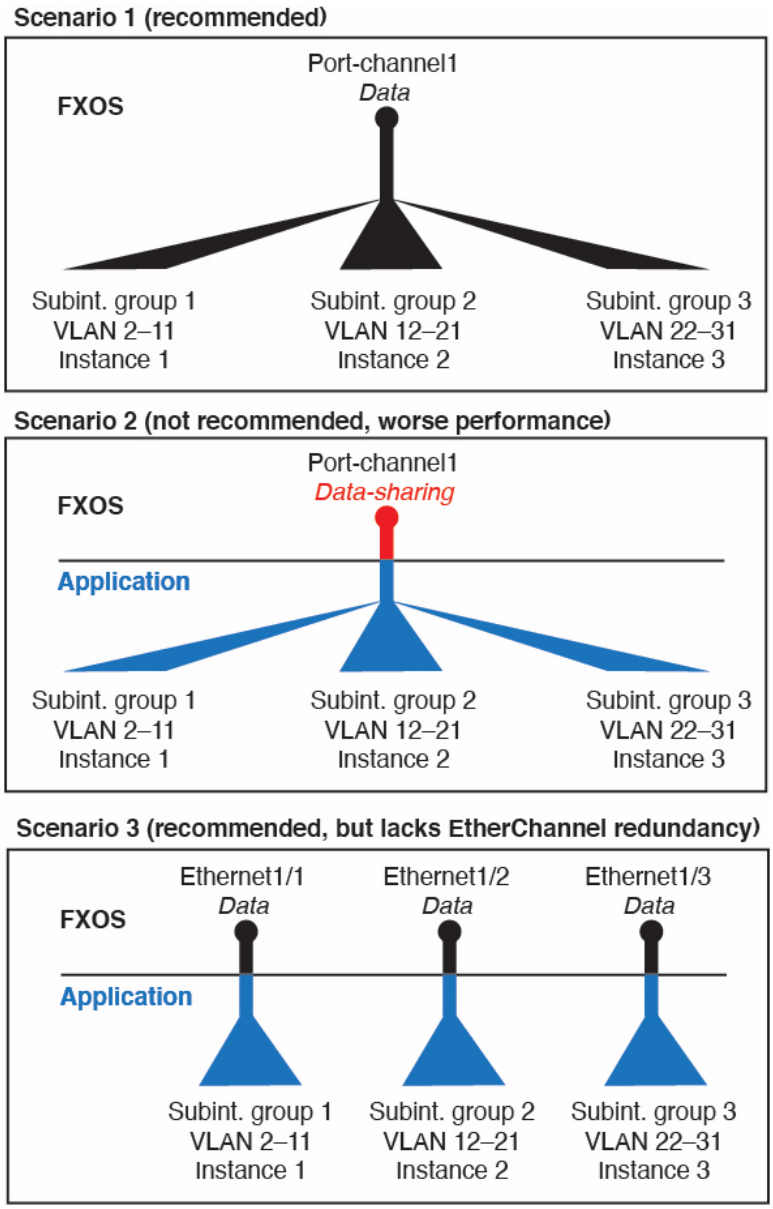

Independent Interface States in the Chassis and in the Application

You can administratively enable and disable interfaces in both the chassis and in the application. For an interface to be operational, the interface must be enabled in both operating systems. Because the interface state is controlled independently, you may have a mismatch between the chassis and application.

The default state of an interface within the application depends on the type of interface. For example, the physical interface or EtherChannel is disabled by default within the application, but a subinterface is enabled by default.

Shared Interface Scalability

Instances can share data-sharing type interfaces. This capability lets you conserve physical interface usage as well as support flexible networking deployments. When you share an interface, the chassis uses unique MAC addresses to forward traffic to the correct instance. However, shared interfaces can cause the forwarding table to grow large due to the need for a full mesh topology within the chassis (every instance must be able to communicate with every other instance that is sharing the same interface). Therefore, there are limits to how many interfaces you can share.

In addition to the forwarding table, the chassis maintains a VLAN group table for VLAN subinterface forwarding. You can create up to 500 VLAN subinterfaces.

See the following limits for shared interface allocation:

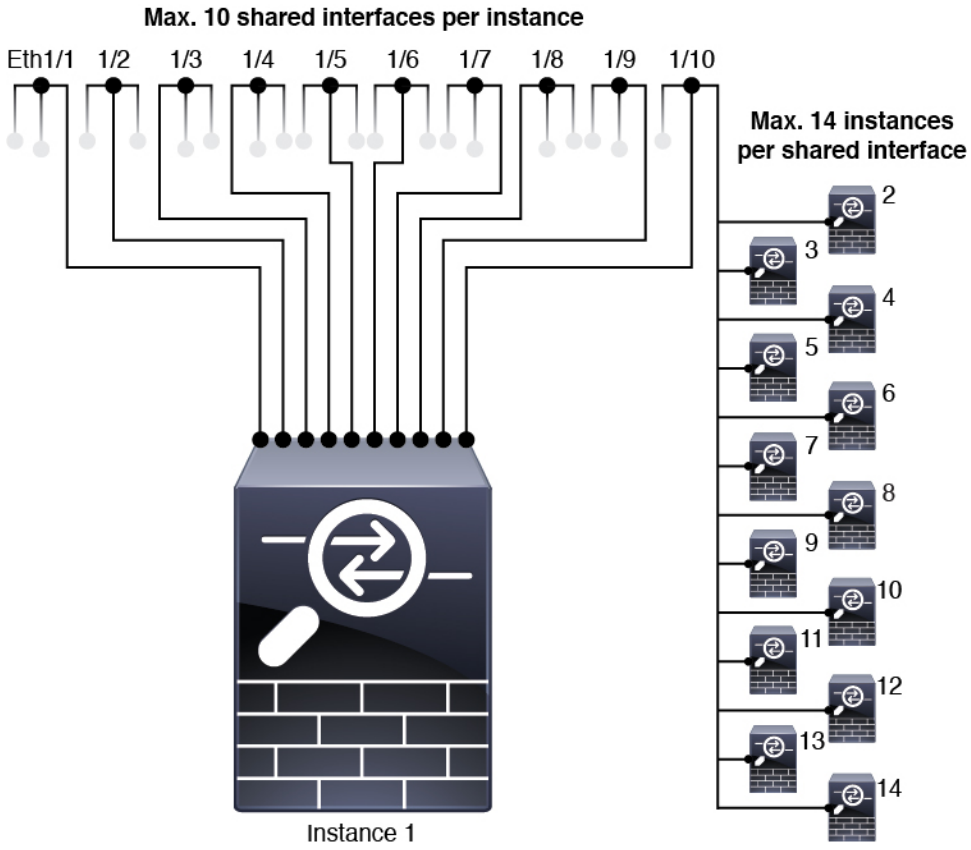

Shared Interface Best Practices

For optimal scalability of the forwarding table, share as few interfaces as possible. Instead, you can create up to 500 VLAN subinterfaces on one or more physical interfaces and then divide the VLANs among the container instances.

When sharing interfaces, follow these practices in the order of most scalable to least scalable:

1. Best—Share subinterfaces under a single parent, and use the same set of subinterfaces with the same group of instances.

For example, create a large EtherChannel to bundle all of your like-kind interfaces together, and then share subinterfaces of that EtherChannel: Port-Channel1.2, 3, and 4 instead of Port-Channel2, Port-Channel3, and Port-Channel4. When you share subinterfaces from a single parent, the VLAN group table provides better scaling of the forwarding table than when sharing physical/EtherChannel interfaces or subinterfaces across parents.

Figure 121: Best: Shared Subinterface Group on One Parent

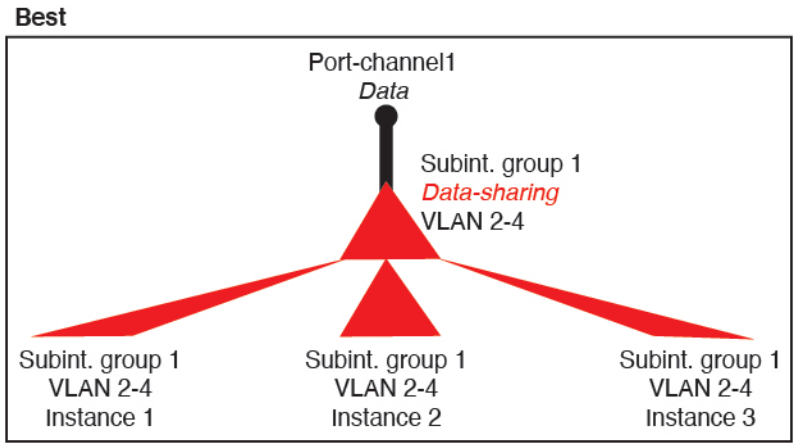

If you do not share the same set of subinterfaces with a group of instances, your configuration can cause more resource usage (more VLAN groups). For example, share Port-Channel1.2, 3, and 4 with instances 1, 2, and 3 (one VLAN group) instead of sharing Port-Channel1.2 and 3 with instances 1 and 2, while sharing Port-Channel1.3 and 4 with instance 3 (two VLAN groups).

Figure 122: Good: Sharing Multiple Subinterface Groups on One Parent

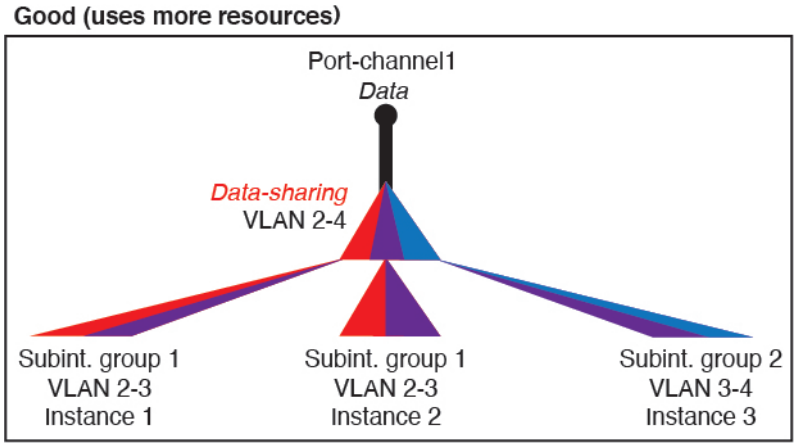

2. Fair—Share subinterfaces across parents.

   For example, share Port-Channel1.2, Port-Channel2.3, and Port-Channel3.4 instead of Port-Channel2, Port-Channel4, and Port-Channel4. Although this usage is not as efficient as only sharing subinterfaces on the same parent, it still takes advantage of VLAN groups.

Figure 123: Fair: Shared Subinterfaces on Separate Parents

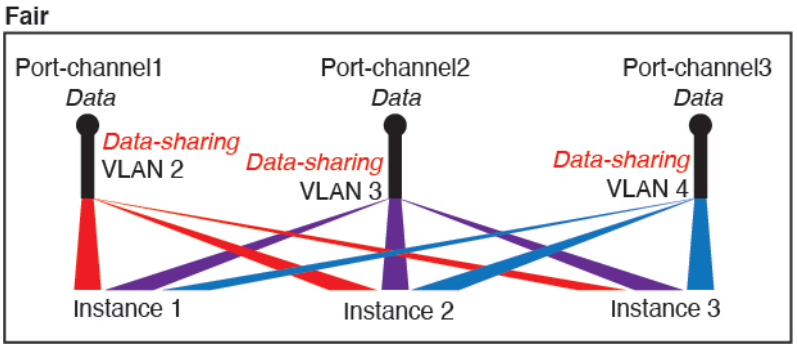

3. Worst—Share individual parent interfaces (physical or EtherChannel).

   This method uses the most forwarding table entries.

Figure 124: Worst: Shared Parent Interfaces

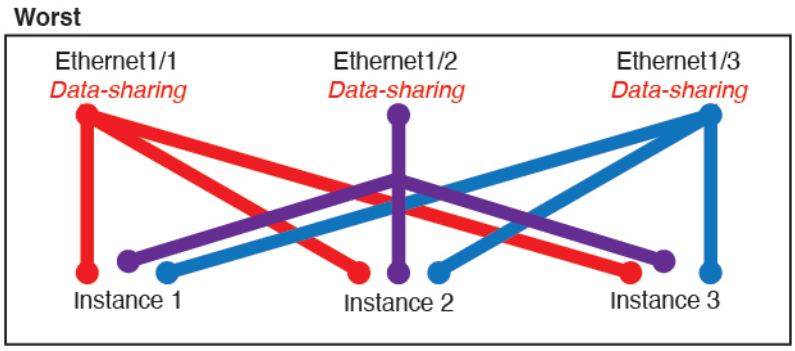

Shared Interface Usage Examples

See the following tables for examples of interface sharing and scalability. The below scenarios assume use of one physical/EtherChannel interface for management shared across all instances, and another physical or EtherChannel interface with dedicated subinterfaces for use with High Availability.

- Table 18: Physical/EtherChannel Interfaces and Instances on a Firepower 9300 with Three SM-44s, on page 305
- Table 19: Subinterfaces on One Parent and Instances on a Firepower 9300 with Three SM-44s, on page 306
- Table 20: Physical/EtherChannel Interfaces and Instances on a Firepower 9300 with One SM-44, on page 308
- Table 21: Subinterfaces on One Parent and Instances on a Firepower 9300 with One SM-44, on page 309

Firepower 9300 with Three SM-44s

The following table applies to three SM-44 security modules on a 9300 using only physical interfaces or EtherChannels. Without subinterfaces, the maximum number of interfaces are limited. Moreover, sharing multiple physical interfaces uses more forwarding table resources than sharing multiple subinterfaces.

Each SM-44 module can support up to 14 instances. Instances are split between modules as necessary to stay within limits.

Table 18: Physical/EtherChannel Interfaces and Instances on a Firepower 9300 with Three SM-44s

Dedicated Interfaces	Shared Interfaces	Number of Instances	% Forwarding Table Used
32: • 8 • 8 • 8 • 8	**0**	**4:** • Instance 1 • Instance 2 • Instance 3 • Instance 4	16%
30: • 15 • 15	**0**	**2:** • Instance 1 • Instance 2	14%
14: • 14 (1 ea.)	**1**	**14:** • Instance 1-Instance 14	46%
33: • 11 (1 ea.) • 11 (1 ea.) • 11 (1 ea.)	**3:** • 1 • 1 • 1	**33:** • Instance 1-Instance 11 • Instance 12-Instance 22 • Instance 23-Instance 33	98%
33: • *11 (1 ea.)* • *11 (1 ea.)* • *12 (1 ea.)*	***3:*** • *1* • *1* • *1*	***34:*** • *Instance 1-Instance 11* • *Instance 12-Instance 22* • *Instance 23-Instance 34*	*102%* DISALLOWED
30: • 30 (1 ea.)	**1**	**6:** • Instance 1-Instance 6	25%
30: • 10 (5 ea.) • 10 (5 ea.) • 10 (5 ea.)	**3:** • 1 • 1 • 1	**6:** • Instance 1-Instance2 • Instance 2-Instance 4 • Instance 5-Instance 6	23%

Dedicated Interfaces	Shared Interfaces	Number of Instances	% Forwarding Table Used
30: • 30 (6 ea.)	**2**	**5:** • Instance 1-Instance 5	28%
30: • 12 (6 ea.) • 18 (6 ea.)	**4:** • 2 • 2	**5:** • Instance 1-Instance2 • Instance 2-Instance 5	26%
24: • 6 • 6 • 6 • 6	**7**	**4:** • Instance 1 • Instance 2 • Instance 3 • Instance 4	44%
24: • 12 (6 ea.) • 12 (6 ea.)	**14:** • 7 • 7	**4:** • Instance 1-Instance2 • Instance 2-Instance 4	41%

The following table applies to three SM-44 security modules on a 9300 using subinterfaces on a single parent physical interface. For example, create a large EtherChannel to bundle all of your like-kind interfaces together, and then share subinterfaces of that EtherChannel. Sharing multiple physical interfaces uses more forwarding table resources than sharing multiple subinterfaces.

Each SM-44 module can support up to 14 instances. Instances are split between modules as necessary to stay within limits.

Table 19: Subinterfaces on One Parent and Instances on a Firepower 9300 with Three SM-44s

Dedicated Subinterfaces	Shared Subinterfaces	Number of Instances	% Forwarding Table Used
168: • 168 (4 ea.)	**0**	**42:** • Instance 1-Instance 42	33%
224: • 224 (16 ea.)	**0**	**14:** • Instance 1-Instance 14	27%
14: • 14 (1 ea.)	**1**	**14:** • Instance 1-Instance 14	46%

Dedicated Subinterfaces	Shared Subinterfaces	Number of Instances	% Forwarding Table Used
33: • 11 (1 ea.) • 11 (1 ea.) • 11 (1 ea.)	**3:** • 1 • 1 • 1	**33:** • Instance 1-Instance 11 • Instance 12-Instance 22 • Instance 23-Instance 33	98%
70: • 70 (5 ea.)	**1**	**14:** • Instance 1-Instance 14	46%
165: • 55 (5 ea.) • 55 (5 ea.) • 55 (5 ea.)	**3:** • 1 • 1 • 1	**33:** • Instance 1-Instance 11 • Instance 12-Instance 22 • Instance 23-Instance 33	98%
70: • 70 (5 ea.)	**2**	**14:** • Instance 1-Instance 14	46%
165: • 55 (5 ea.) • 55 (5 ea.) • 55 (5 ea.)	**6:** • 2 • 2 • 2	**33:** • Instance 1-Instance 11 • Instance 12-Instance 22 • Instance 23-Instance 33	98%
70: • 70 (5 ea.)	**10**	**14:** • Instance 1-Instance 14	46%
165: • *55 (5 ea.)* • *55 (5 ea.)* • *55 (5 ea.)*	***30:*** • *10* • *10* • *10*	***33:*** • *Instance 1-Instance 11* • *Instance 12-Instance 22* • *Instance 23-Instance 33*	*102%* DISALLOWED

Firepower 9300 with One SM-44

The following table applies to the Firepower 9300 with one SM-44 using only physical interfaces or EtherChannels. Without subinterfaces, the maximum number of interfaces are limited. Moreover, sharing multiple physical interfaces uses more forwarding table resources than sharing multiple subinterfaces.

The Firepower 9300 with one SM-44 can support up to 14 instances.

Table 20: Physical/EtherChannel Interfaces and Instances on a Firepower 9300 with One SM-44

Dedicated Interfaces	Shared Interfaces	Number of Instances	% Forwarding Table Used
32: • 8 • 8 • 8 • 8	**0**	**4:** • Instance 1 • Instance 2 • Instance 3 • Instance 4	16%
30: • 15 • 15	**0**	**2:** • Instance 1 • Instance 2	14%
14: • 14 (1 ea.)	**1**	**14:** • Instance 1-Instance 14	46%
14: • 7 (1 ea.) • 7 (1 ea.)	**2:** • **1** • **1**	**14:** • Instance 1-Instance 7 • Instance 8-Instance 14	37%
32: • 8 • 8 • 8 • 8	**1**	**4:** • Instance 1 • Instance 2 • Instance 3 • Instance 4	21%
32: • 16 (8 ea.) • 16 (8 ea.)	**2**	**4:** • Instance 1-Instance 2 • Instance 3-Instance 4	20%
32: • 8 • 8 • 8 • 8	**2**	**4:** • Instance 1 • Instance 2 • Instance 3 • Instance 4	25%

Dedicated Interfaces	Shared Interfaces	Number of Instances	% Forwarding Table Used
32: • 16 (8 ea.) • 16 (8 ea.)	**4:** • 2 • 2	**4:** • Instance 1-Instance 2 • Instance 3-Instance 4	24%
24: • 8 • 8 • 8	**8**	**3:** • Instance 1 • Instance 2 • Instance 3	37%
10: • 10 (2 ea.)	**10**	**5:** • Instance 1-Instance 5	69%
10: • 6 (2 ea.) • 4 (2 ea.)	**20:** • 10 • 10	**5:** • Instance 1-Instance 3 • Instance 4-Instance 5	59%
14: • *12 (2 ea.)*	***10***	***7:*** • *Instance 1-Instance 7*	*109%* DISALLOWED

The following table applies to the Firepower 9300 with one SM-44 using subinterfaces on a single parent physical interface. For example, create a large EtherChannel to bundle all of your like-kind interfaces together, and then share subinterfaces of that EtherChannel. Sharing multiple physical interfaces uses more forwarding table resources than sharing multiple subinterfaces.

The Firepower 9300 with one SM-44 can support up to 14 instances.

Table 21: Subinterfaces on One Parent and Instances on a Firepower 9300 with One SM-44

Dedicated Subinterfaces	Shared Subinterfaces	Number of Instances	% Forwarding Table Used
112: • 112 (8 ea.)	**0**	**14:** • Instance 1-Instance 14	17%
224: • 224 (16 ea.)	**0**	**14:** • Instance 1-Instance 14	17%
14: • 14 (1 ea.)	**1**	**14:** • Instance 1-Instance 14	46%

Dedicated Subinterfaces	Shared Subinterfaces	Number of Instances	% Forwarding Table Used
14: • 7 (1 ea.) • 7 (1 ea.)	**2:** • 1 • 1	**14:** • Instance 1-Instance 7 • Instance 8-Instance 14	37%
112: • 112 (8 ea.)	**1**	**14:** • Instance 1-Instance 14	46%
112: • 56 (8 ea.) • 56 (8 ea.)	**2:** • 1 • 1	**14:** • Instance 1-Instance 7 • Instance 8-Instance 14	37%
112: • 112 (8 ea.)	**2**	**14:** • Instance 1-Instance 14	46%
112: • 56 (8 ea.) • 56 (8 ea.)	**4:** • 2 • 2	**14:** • Instance 1-Instance 7 • Instance 8-Instance 14	37%
140: • 140 (10 ea.)	**10**	**14:** • Instance 1-Instance 14	46%
140: • 70 (10 ea.) • 70 (10 ea.)	**20:** • 10 • 10	**14:** • Instance 1-Instance 7 • Instance 8-Instance 14	37%

Viewing Shared Interface Resources

To view forwarding table and VLAN group usage, see the **Devices & Network** > **Interface Forwarding Utilization** area. For example:

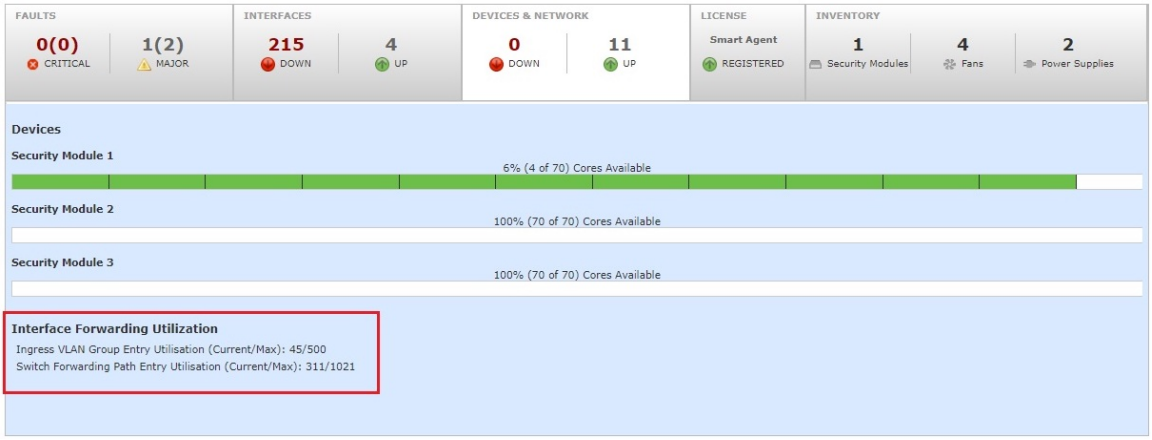

Inline Set Link State Propagation for the Firewall Threat Defense

An inline set acts like a bump on the wire, and binds two interfaces together to slot into an existing network. This function allows the system to be installed in any network environment without the configuration of adjacent network devices. Inline interfaces receive all traffic unconditionally, but all traffic received on these interfaces is retransmitted out of an inline set unless explicitly dropped.

When you configure an inline set in the Firewall Threat Defense application and enable link state propagation, the Firewall Threat Defense sends inline set membership to the FXOS chassis. Link state propagation means that the chassis automatically brings down the second interface in the inline interface pair when one of the interfaces in an inline set goes down. When the downed interface comes back up, the second interface automatically comes back up, also. In other words, if the link state of one interface changes, the chassis senses the change and updates the link state of the other interface to match it. Note that the chassis requires up to 4 seconds to propagate link state changes. Link state propagation is especially useful in resilient network environments where routers are configured to reroute traffic automatically around network devices that are in a failure state.

Note Do not enable Hardware Bypass and link state propagation for the same inline set.

About Logical Devices

A logical device lets you run one application instance (either ASA or Firewall Threat Defense) and also one optional decorator application (Radware DefensePro) to form a service chain.

When you add a logical device, you also define the application instance type and version, assign interfaces, and configure bootstrap settings that are pushed to the application configuration.

Note For the Firepower 9300, you can install different application types (ASA and Firewall Threat Defense) on separate modules in the chassis. You can also run different versions of an application instance type on separate modules.

Standalone and Clustered Logical Devices

You can add the following logical device types:

- Standalone—A standalone logical device operates as a standalone unit or as a unit in a High Availability pair.
- Cluster—A clustered logical device lets you group multiple units together, providing all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices. Multiple module devices, like the Firepower 9300, support intra-chassis clustering. For the Firepower 9300, all three modules must participate in the cluster, for both native and container instances. The Firewall Device Manager does not support clustering.

Logical Device Application Instances: Container and Native

Application instances run in the following deployment types:

- Native instance—A native instance uses all of the resources (CPU, RAM, and disk space) of the security module/engine, so you can only install one native instance.
- Container instance—A container instance uses a subset of resources of the security module/engine, so you can install multiple container instances. Multi-instance capability is only supported for the Firewall Threat Defense using Firewall Management Center; it is not supported for the ASA or the Firewall Threat Defense using Firewall Device Manager.

Note Multi-instance capability is similar to ASA multiple context mode, although the implementation is different. Multiple context mode partitions a single application instance, while multi-instance capability allows independent container instances. Container instances allow hard resource separation, separate configuration management, separate reloads, separate software updates, and full Firewall Threat Defense feature support. Multiple context mode, due to shared resources, supports more contexts on a given platform. Multiple context mode is not available on the Firewall Threat Defense.

For the Firepower 9300, you can use a native instance on some modules, and container instances on the other module(s).

Container Instance Interfaces

To provide flexible physical interface use for container instances, you can create VLAN subinterfaces in FXOS and also share interfaces (VLAN or physical) between multiple instances. Native instances cannot use VLAN subinterfaces or shared interfaces. A multi-instance cluster cannot use VLAN subinterfaces or shared interfaces. An exception is made for the cluster control link, which can use a subinterface of the Cluster EtherChannel. See Shared Interface Scalability, on page 302 and Add a VLAN Subinterface for Container Instances, on page 335.

Note This document discusses *FXOS* VLAN subinterfaces only. You can separately create subinterfaces within the Firewall Threat Defense application. See FXOS Interfaces vs. Application Interfaces, on page 300 for more information.

How the Chassis Classifies Packets

Each packet that enters the chassis must be classified, so that the chassis can determine to which instance to send a packet.

- Unique Interfaces—If only one instance is associated with the ingress interface, the chassis classifies the packet into that instance. For bridge group member interfaces (in transparent mode or routed mode), inline sets, or passive interfaces, this method is used to classify packets at all times.
- Unique MAC Addresses—The chassis automatically generates unique MAC addresses for all interfaces, including shared interfaces. If multiple instances share an interface, then the classifier uses unique MAC addresses assigned to the interface in each instance. An upstream router cannot route directly to an instance without unique MAC addresses. You can also set the MAC addresses manually when you configure each interface within the application.

Note If the destination MAC address is a multicast or broadcast MAC address, the packet is duplicated and delivered to each instance.

Classification Examples

Packet Classification with a Shared Interface Using MAC Addresses

The following figure shows multiple instances sharing an outside interface. The classifier assigns the packet to Instance C because Instance C includes the MAC address to which the router sends the packet.

Figure 125: Packet Classification with a Shared Interface Using MAC Addresses

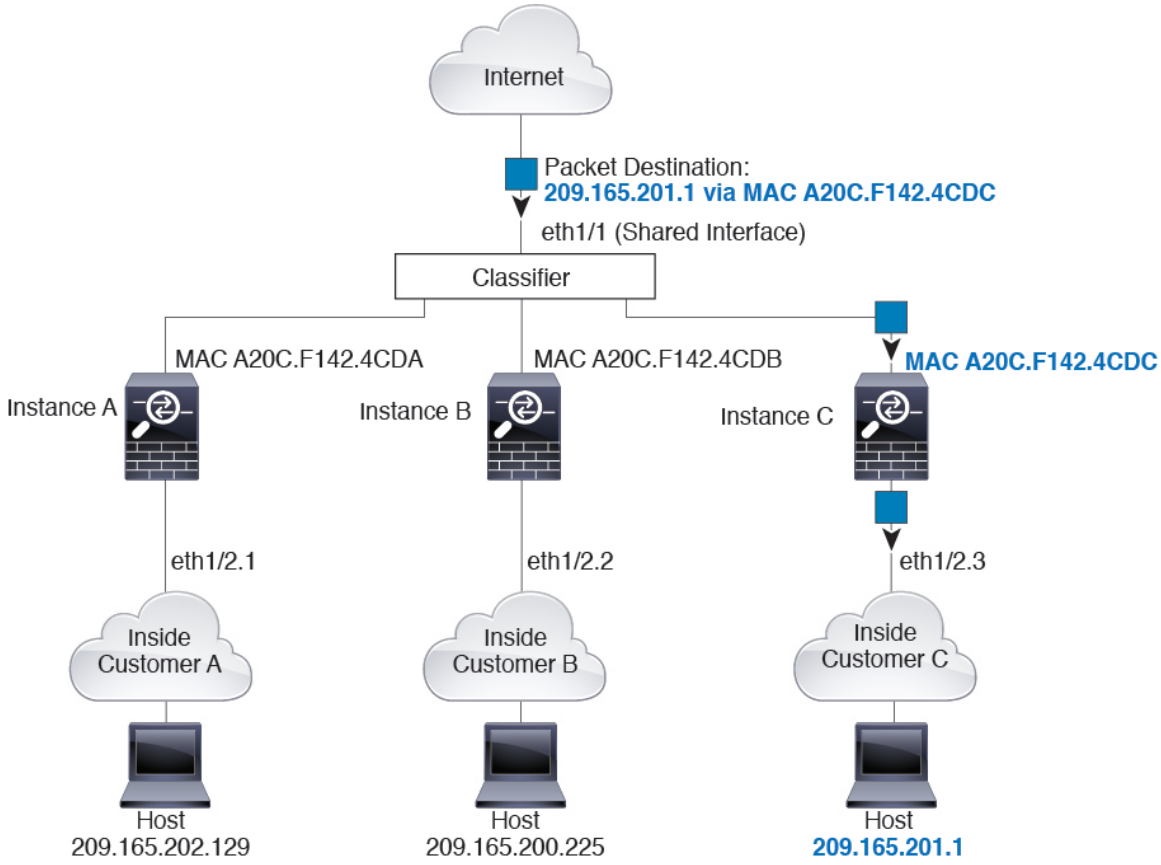

Incoming Traffic from Inside Networks

Note that all new incoming traffic must be classified, even from inside networks. The following figure shows a host on the Instance C inside network accessing the internet. The classifier assigns the packet to Instance C because the ingress interface is Ethernet 1/2.3, which is assigned to Instance C.

Figure 126: Incoming Traffic from Inside Networks

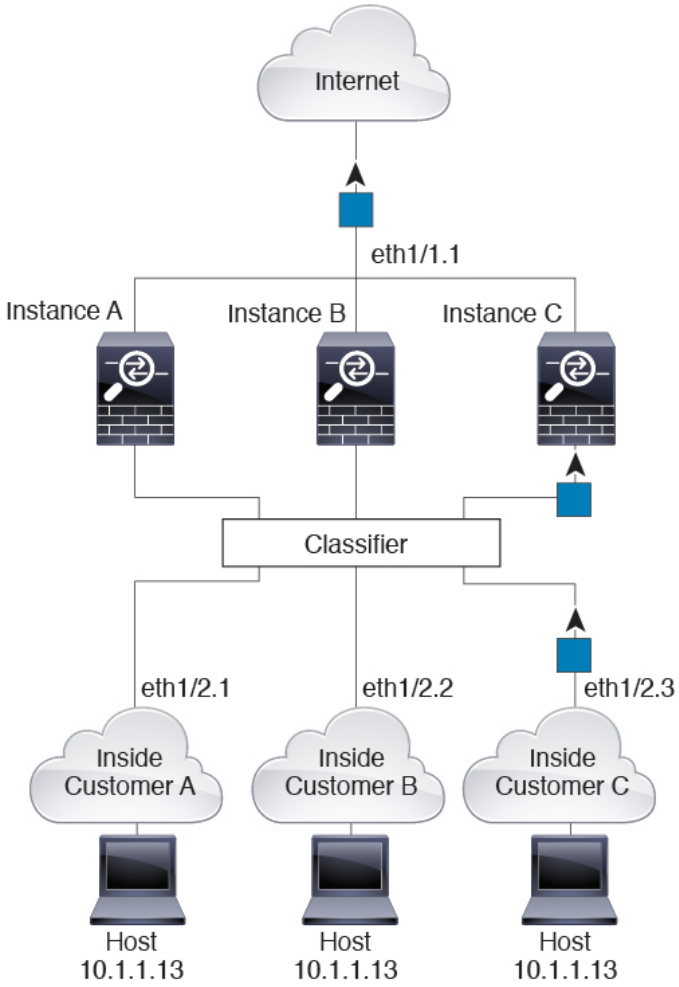

Transparent Firewall Instances

For transparent firewalls, you must use unique interfaces. The following figure shows a packet destined to a host on the Instance C inside network from the internet. The classifier assigns the packet to Instance C because the ingress interface is Ethernet 1/2.3, which is assigned to Instance C.

Figure 127: Transparent Firewall Instances

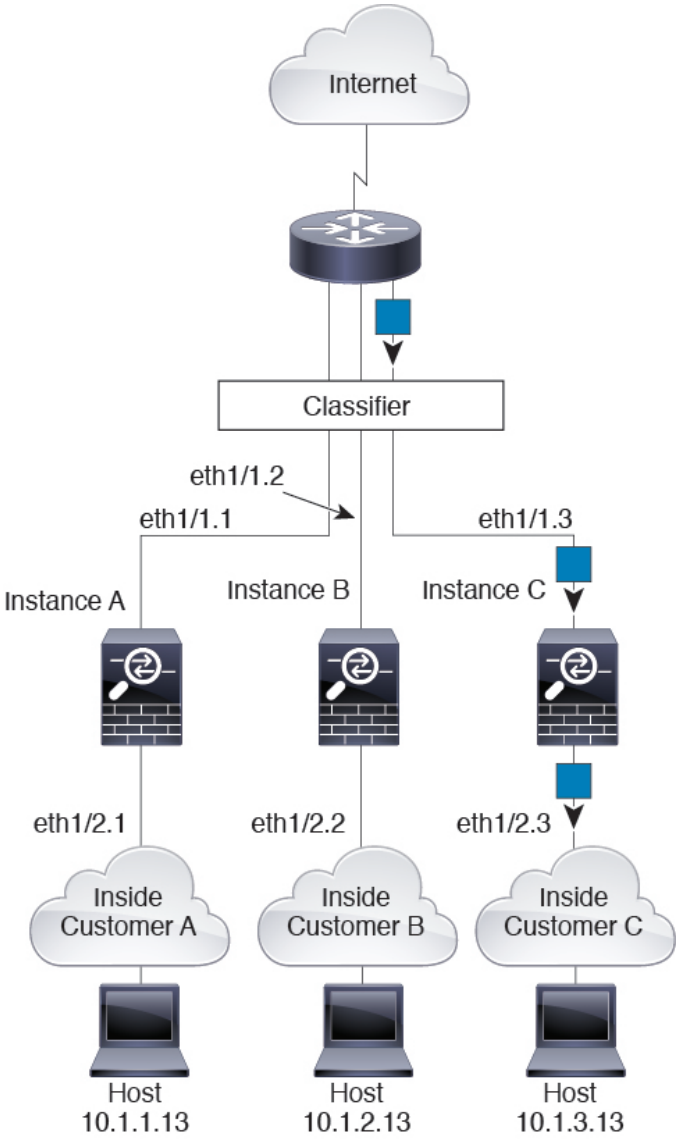

Inline Sets

For inline sets, you must use unique interfaces and they must be physical interfaces or EtherChannels. The following figure shows a packet destined to a host on the Instance C inside network from the internet. The classifier assigns the packet to Instance C because the ingress interface is Ethernet 1/5, which is assigned to Instance C.

Figure 128: Inline Sets

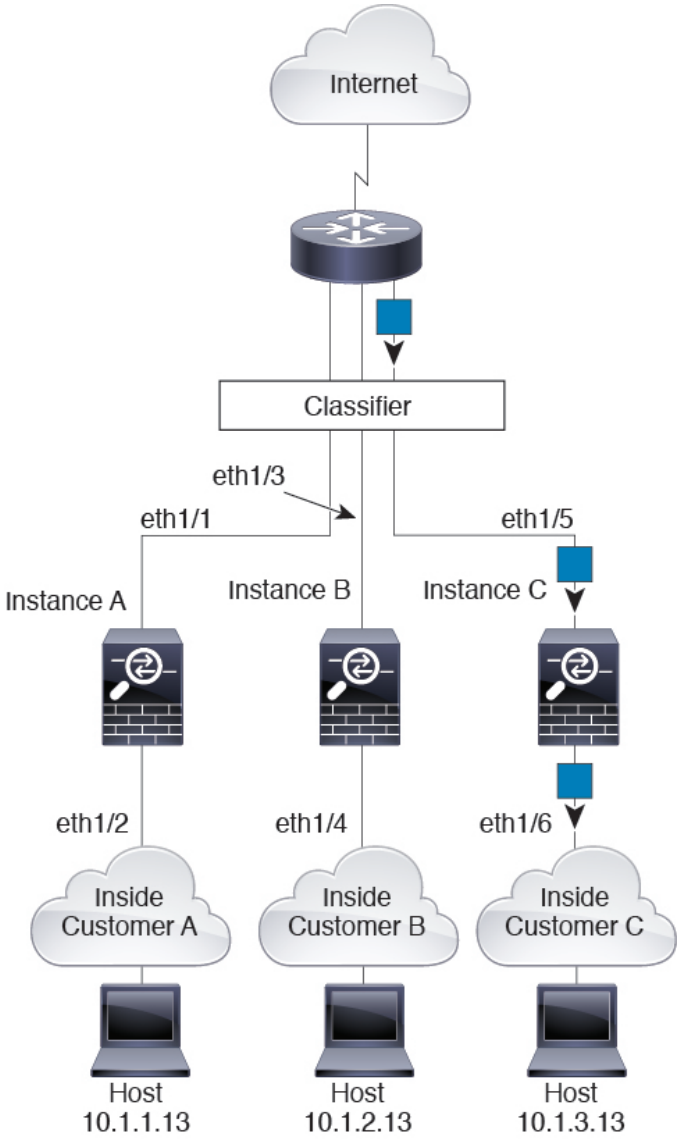

Cascading Container Instances

Placing an instance directly in front of another instance is called *cascading instances*; the outside interface of one instance is the same interface as the inside interface of another instance. You might want to cascade instances if you want to simplify the configuration of some instances by configuring shared parameters in the top instance.

The following figure shows a gateway instance with two instances behind the gateway.

Figure 129: Cascading Instances

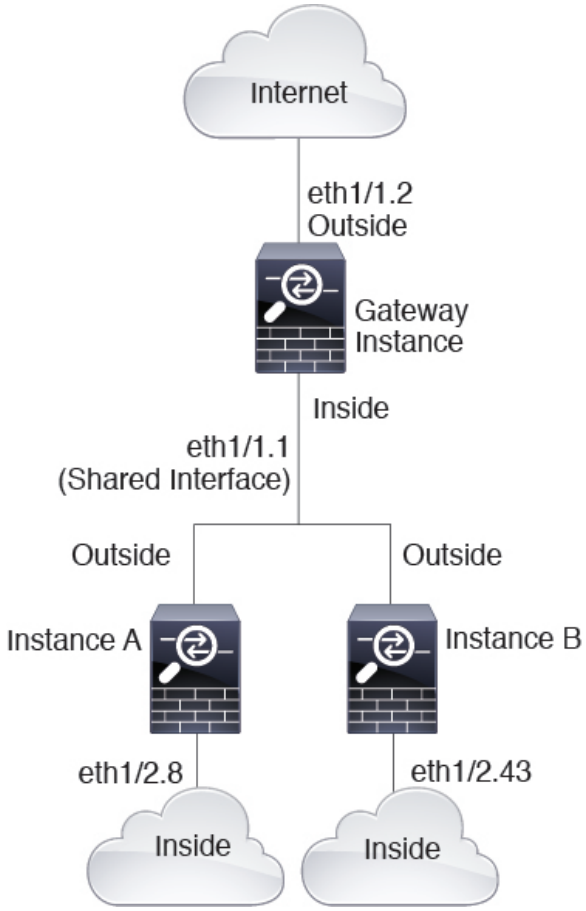

Note Do not use cascading instances (using a shared interface) with High Availability. After a failover occurs and the standby unit rejoins, MAC addresses can overlap temporarily and cause an outage. You should instead use unique interfaces for the gateway instance and inside instance using an external switch to pass traffic between the instances.

Typical Multi-Instance Deployment

The following example includes three container instances in routed firewall mode. They include the following interfaces:

- Management—All instances use the Port-Channel1 interface (management type). This EtherChannel includes two 10 Gigibit Ethernet interfaces. Within each application, the interface uses a unique IP address on the same management network.
- Inside—Each instance uses a subinterface on Port-Channel2 (data type). This EtherChannel includes two 10 Gigibit Ethernet interfaces. Each subinterface is on a separate network.

- Outside—All instances use the Port-Channel3 interface (data-sharing type). This EtherChannel includes two 10 Gigibit Ethernet interfaces. Within each application, the interface uses a unique IP address on the same outside network.
- Failover—Each instance uses a subinterface on Port-Channel4 (data type). This EtherChannel includes two 10 Gigibit Ethernet interfaces. Each subinterface is on a separate network.

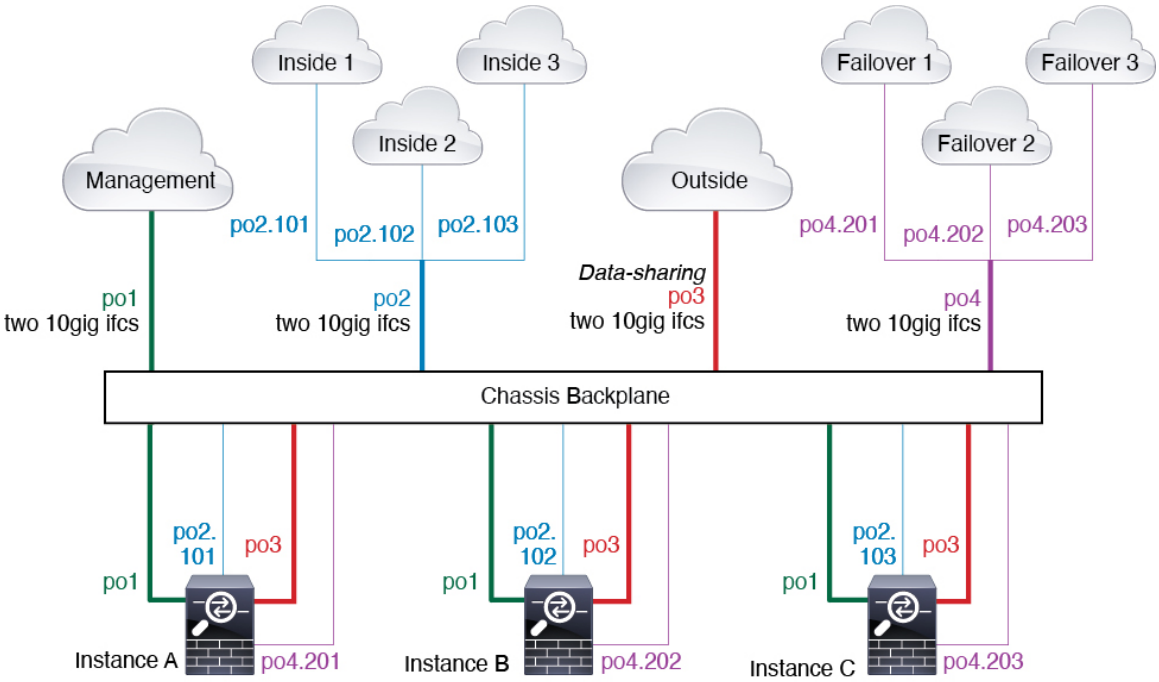

Automatic MAC Addresses for Container Instance Interfaces

The chassis automatically generates MAC addresses for instance interfaces, and guarantees that a shared interface in each instance uses a unique MAC address.

If you manually assign a MAC address to a shared interface within the instance, then the manually-assigned MAC address is used. If you later remove the manual MAC address, the autogenerated address is used. In the rare circumstance that the generated MAC address conflicts with another private MAC address in your network, we suggest that you manually set the MAC address for the interface within the instance.

Because autogenerated addresses start with A2, you should not start manual MAC addresses with A2 due to the risk of overlapping addresses.

The chassis generates the MAC address using the following format:

A2*xx.yyzz.zzzz*

Where *xx.yy* is a user-defined prefix or a system-defined prefix, and *zz.zzzz* is an internal counter generated by the chassis. The system-defined prefix matches the lower 2 bytes of the first MAC address in the burned-in MAC address pool that is programmed into the IDPROM. Use **connect fxos**, then **show module** to view the MAC address pool. For example, if the range of MAC addresses shown for module 1 is b0aa.772f.f0b0 to b0aa.772f.f0bf, then the system prefix will be f0b0.

The user-defined prefix is an integer that is converted into hexadecimal. For an example of how the user-defined prefix is used, if you set a prefix of 77, then the chassis converts 77 into the hexadecimal value 004D (*yyxx*). When used in the MAC address, the prefix is reversed (*xxyy*) to match the chassis native form:

A2**4D.00***zz.zzzz*

For a prefix of 1009 (03F1), the MAC address is:

A2**F1.03***zz.zzzz*

Container Instance Resource Management

To specify resource usage per container instance, create one or more resource profiles in FXOS. When you deploy the logical device/application instance, you specify the resource profile that you want to use. The resource profile sets the number of CPU cores; RAM is dynamically allocated according to the number of cores, and disk space is set to 40 GB per instance. To view the available resources per model, see Requirements and Prerequisites for Container Instances, on page 323. To add a resource profile, see Add a Resource Profile for Container Instances, on page 337.

Performance Scaling Factor for Multi-Instance Capability

The maximum throughput (connections, VPN sessions, and TLS proxy sessions) for a platform is calculated for a native instance's use of memory and CPU (and this value is shown in **show resource usage**). If you use multiple instances, then you need to calculate the throughput based on the percentage of CPU cores that you assign to the instance. For example, if you use a container instance with 50% of the cores, then you should initially calculate 50% of the throughput. Moreover, the throughput available to a container instance may be less than that available to a native instance.

For detailed instructions on calculating the throughput for instances, see https://www.cisco.com/c/en/us/products/collateral/security/firewalls/white-paper-c11-744750.html.

Container Instances and High Availability

You can use High Availability using a container instance on 2 separate chassis; for example, if you have 2 chassis, each with 10 instances, you can create 10 High Availability pairs. Note that High Availability is not configured in FXOS; configure each High Availability pair in the application manager.

For detailed requirements, see Requirements and Prerequisites for High Availability, on page 324 and Add a High Availability Pair, on page 343.

Container Instances and Clustering

You can create a cluster of container instances using one container instance per security module/engine. See Requirements and Prerequisites for Clustering, on page 325 for detailed requirements.

Licenses for Container Instances

All licenses are consumed per security engine/chassis (for the Firepower 4100) or per security module (for the Firepower 9300), and not per container instance. See the following details:

- Essentials licenses are automatically assigned: one per security module/engine.
- Feature licenses are manually assigned to each instance; but you only consume one license per feature per security module/engine. For example, for the Firepower 9300 with 3 security modules, you only need

one URL Filtering license per module for a total of 3 licenses, regardless of the number of instances in use.

For example:

Table 22: Sample License Usage for Container Instances on a Firepower 9300

Firepower 9300	Instance	Licenses
Security Module 1	Instance 1	Essentials, URL Filtering, Malware Defense
	Instance 2	Essentials, URL Filtering
	Instance 3	Essentials, URL Filtering
Security Module 2	Instance 4	Essentials, IPS
	Instance 5	Essentials, URL Filtering, Malware Defense, IPS
Security Module 3	Instance 6	Essentials, Malware Defense, IPS
	Instance 7	Essentials, IPS

Table 23: Total Number of Licenses

Essentials	URL Filtering	Malware Defense	IPS
3	2	3	2

Requirements and Prerequisites for Logical Devices

See the following sections for requirements and prerequisites.

Requirements and Prerequisites for Hardware and Software Combinations

The Firepower 4100/9300 supports multiple models, security modules, application types, and high availability and scalability features. See the following requirements for allowed combinations.

Firepower 9300 Requirements

The Firepower 9300 includes 3 security module slots and multiple types of security modules. See the following requirements:

- Security Module Types—You can install modules of different types in the Firepower 9300. For example, you can install the SM-48 as module 1, SM-40 as module 2, and SM-56 as module 3.
- Native instance Clustering—All security modules in the cluster, whether it is intra-chassis or inter-chassis, must be the same type. You can have different quantities of installed security modules in each chassis, although all modules present in the chassis must belong to the cluster including any empty slots. For

example, you can install 2 SM-40s in chassis 1, and 3 SM-40s in chassis 2. You cannot use clustering if you install 1 SM-48 and 2 SM-40s in the same chassis.

- Container instance Clustering—You can create a cluster using instances on different model types. For example, you can create a cluster using an instance on a Firepower 9300 SM-56, SM-48, and SM-40. You *cannot* mix the Firepower 9300 and the Firepower 4100 in the same cluster, however.

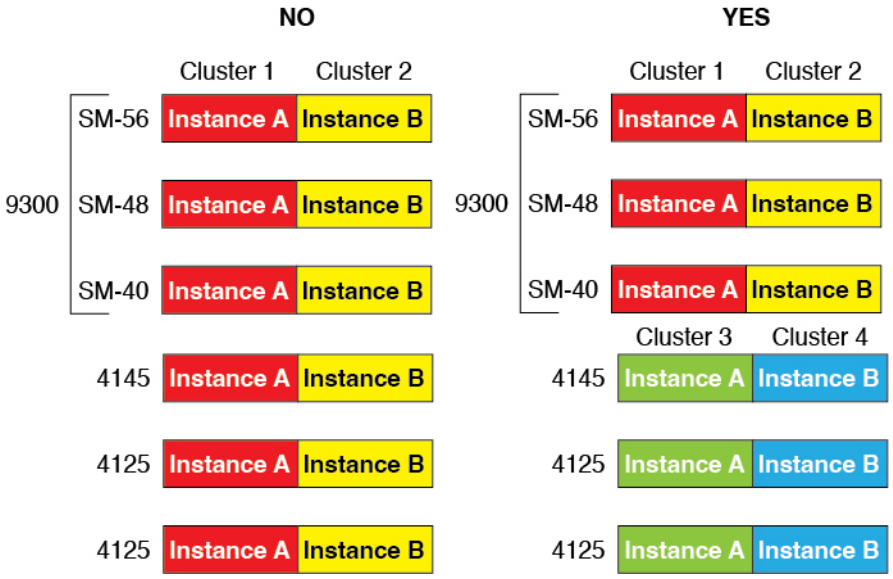

- High Availability—High Availability is only supported between same-type modules on the Firepower 9300. However, the two chassis can include mixed modules. For example, each chassis has an SM-40, SM-48, and SM-56. You can create High Availability pairs between the SM-40 modules, between the SM-48 modules, and between the SM-56 modules.
- ASA and Firewall Threat Defense application types—You can install different application types on separate modules in the chassis. For example, you can install ASA on module 1 and module 2, and Firewall Threat Defense on module 3.
- ASA or Firewall Threat Defense versions—You can run different versions of an application instance type on separate modules, or as separate container instances on the same module. For example, you can install the Firewall Threat Defense 6.3 on module 1, Firewall Threat Defense 6.4 on module 2, and Firewall Threat Defense 6.5 on module 3.

Firepower 4100 Requirements

The Firepower 4100 comes in multiple models. See the following requirements:

- Native and Container instances—When you install a container instance on a Firepower 4100, that device can only support other container instances. A native instance uses all of the resources for a device, so you can only install a single native instance on the device.
- Native instance Clustering—All chassis in the cluster must be the same model.
- Container instance Clustering—You can create a cluster using instances on different model types. For example, you can create a cluster using an instance on a Firepower 4145 and a 4125. You *cannot* mix the Firepower 9300 and the Firepower 4100 in the same cluster, however.

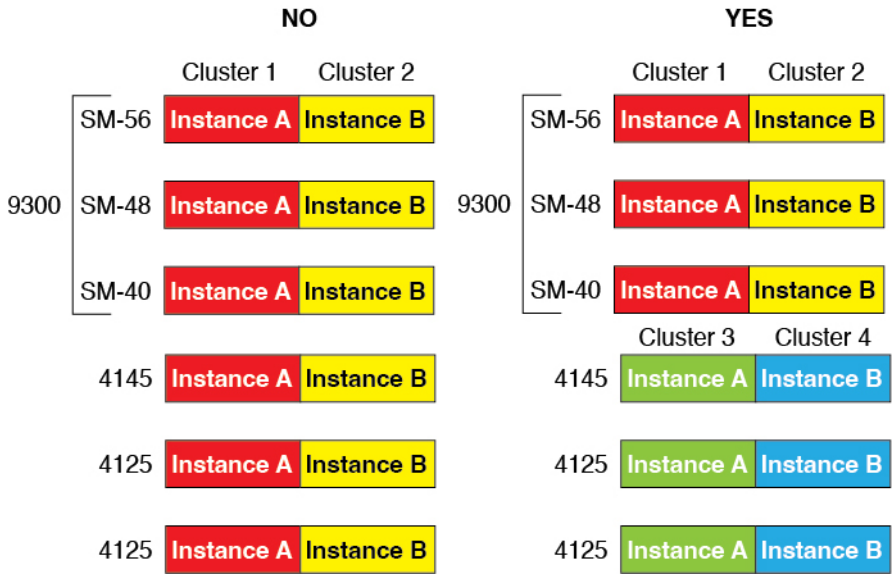

- High Availability—High Availability is only supported between same-type models.
- ASA and Firewall Threat Defense application types—The Firepower 4100 can only run a single application type.
- The Firewall Threat Defense container instance versions—You can run different versions of Firewall Threat Defense as separate container instances on the same module.

Requirements and Prerequisites for Container Instances

For information about high-availability or clustering requirements with multi-instance, see Requirements and Prerequisites for High Availability, on page 324 and see Requirements and Prerequisites for Clustering, on page 325.

Supported Application Types

- The Firewall Threat Defense using Firewall Management Center

Maximum Container Instances and Resources per Model

For each container instance, you can specify the number of CPU cores to assign to the instance. RAM is dynamically allocated according to the number of cores, and disk space is set to 40 GB per instance.

Table 24: Maximum Container Instances and Resources per Model

Model	Max. Container Instances	Available CPU Cores	Available RAM	Available Disk Space
Firepower 4112	3	22	78 GB	308 GB
Firepower 4115	7	46	162 GB	308 GB
Firepower 4125	10	62	162 GB	644 GB

Model	Max. Container Instances	Available CPU Cores	Available RAM	Available Disk Space
Firepower 4140	7	70	222 GB	311.8 GB
Firepower 4145	14	86	344 GB	608 GB
Firepower 9300 SM-40 security module	13	78	334 GB	1359 GB
Firepower 9300 SM-48 security module	15	94	334 GB	1341 GB
Firepower 9300 SM-56 security module	18	110	334 GB	1314 GB

Firewall Management Center Requirements

For all instances on a Firepower 4100 chassis or Firepower 9300 module, you must use the same Firewall Management Center due to the licensing implementation.

Requirements and Prerequisites for High Availability

- The two units in a High Availability Failover configuration must:
  - Be on a separate chassis; intra-chassis High Availability for the Firepower 9300 is not supported.
  - Be the same model.
  - Have the same interfaces assigned to the High Availability logical devices.
  - Have the same number and types of interfaces. All interfaces must be preconfigured in FXOS identically before you enable High Availability.
- High Availability is only supported between same-type modules on the Firepower 9300; but the two chassis can include mixed modules. For example, each chassis has an SM-56, SM-48, and SM-40. You can create High Availability pairs between the SM-56 modules, between the SM-48 modules, and between the SM-40 modules.
- For container instances, each unit must use the same resource profile attributes.
- For container instances: Do not use cascading instances (using a shared interface) with High Availability. After a failover occurs and the standby unit rejoins, MAC addresses can overlap temporarily and cause an outage. You should instead use unique interfaces for the gateway instance and inside instance using an external switch to pass traffic between the instances.
- For other High Availability system requirements, see High Availability System Requirements, on page 423.

Requirements and Prerequisites for Clustering

Cluster Model Support

The Threat Defense supports clustering on the following models:

- Firepower 9300— You can include up to 16 nodes in the cluster. For example, you can use 1 module in 16 chassis, or 2 modules in 8 chassis, or any combination that provides a maximum of 16 modules. Supports clustering with multiple chassis and clustering isolated to security modules within one chassis.
- Firepower 4100—Supported for up to 16 nodes using clustering with multiple chassis.

User Roles

- Admin
- Access Admin
- Network Admin

Clustering Hardware and Software Requirements

All chassis in a cluster:

- Native instance clustering—For the Firepower 4100: All chassis must be the same model. For the Firepower 9300: All security modules must be the same type. For example, if you use clustering, all modules in the Firepower 9300 must be SM-40s. You can have different quantities of installed security modules in each chassis, although all modules present in the chassis must belong to the cluster including any empty slots.
- Container instance clustering—We recommend that you use the same security module or chassis model for each cluster instance. However, you can mix and match container instances on different Firepower 9300 security module types or Firepower 4100 models in the same cluster if required. You cannot mix Firepower 9300 and 4100 instances in the same cluster. For example, you can create a cluster using an instance on a Firepower 9300 SM-56, SM-48, and SM-40. Or you can create a cluster on a Firepower 4145 and a 4125.

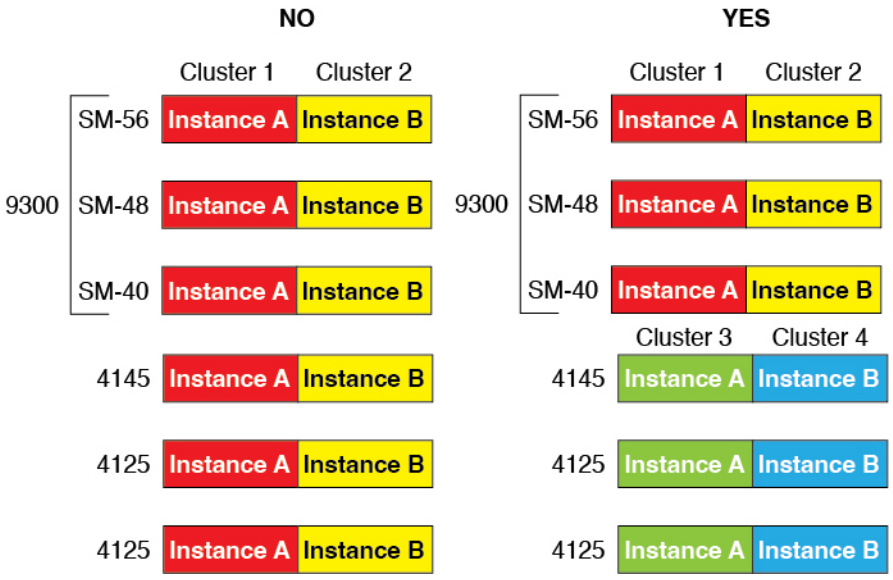

- Must run the identical FXOS and application software except at the time of an image upgrade. Mismatched software versions can lead to poor performance, so be sure to upgrade all nodes in the same maintenance window.
- Must include the same interface configuration for interfaces you assign to the cluster, such as the same Management interface, EtherChannels, active interfaces, speed and duplex, and so on. You can use different network module types on the chassis as long as the capacity matches for the same interface IDs and interfaces can successfully bundle in the same spanned EtherChannel. Note that all data interfaces must be EtherChannels in clusters with multiple chassis. If you change the interfaces in FXOS after you enable clustering (by adding or removing interface modules, or configuring EtherChannels, for example), then perform the same changes on each chassis, starting with the data nodes, and ending with the control node.
- Must use the same NTP server. For Firewall Threat Defense, the Firewall Management Center must also use the same NTP server. Do not set the time manually.

Multi-Instance Clustering Requirements

- No intra-security-module/engine clustering—For a given cluster, you can only use a single container instance per security module/engine. You cannot add 2 container instances to the same cluster if they are running on the same module.

- Mix and match clusters and standalone instances—Not all container instances on a security module/engine need to belong to a cluster. You can use some instances as standalone or High Availability nodes. You can also create multiple clusters using separate instances on the same security module/engine.

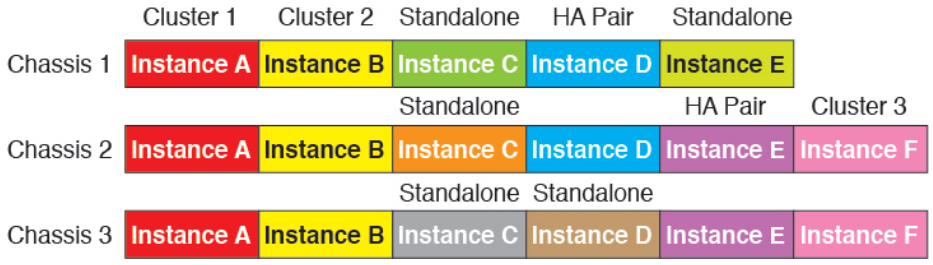

- All 3 modules in a Firepower 9300 must belong to the cluster—For the Firepower 9300, a cluster requires a single container instance on all 3 modules. You cannot create a cluster using instances on module 1 and 2, and then use a native instance on module 3, or example.

- Match resource profiles—We recommend that each node in the cluster use the same resource profile attributes; however, mismatched resources are allowed when changing cluster nodes to a different resource profile, or when using different models.
- Dedicated cluster control link—For clusters with multiple chassis, each cluster needs a dedicated cluster control link. For example, each cluster can use a separate subinterface on the same cluster-type EtherChannel, or use separate EtherChannels.

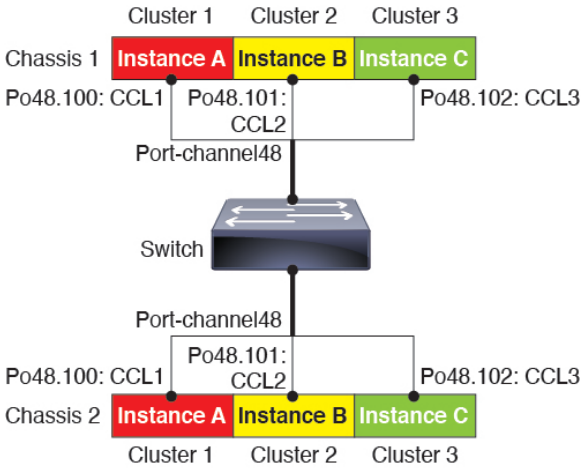

- No shared interfaces—Shared-type interfaces are not supported with clustering. However, the same Management and Eventing interfaces can used by multiple clusters.
- No subinterfaces—A multi-instance cluster cannot use FXOS-defined VLAN subinterfaces. An exception is made for the cluster control link, which can use a subinterface of the Cluster EtherChannel.
- Mix chassis models—We recommend that you use the same security module or chassis model for each cluster instance. However, you can mix and match container instances on different Firepower 9300

security module types or Firepower 4100 models in the same cluster if required. You cannot mix Firepower 9300 and 4100 instances in the same cluster. For example, you can create a cluster using an instance on a Firepower 9300 SM-56, SM-48, and SM-40. Or you can create a cluster on a Firepower 4145 and a 4125.

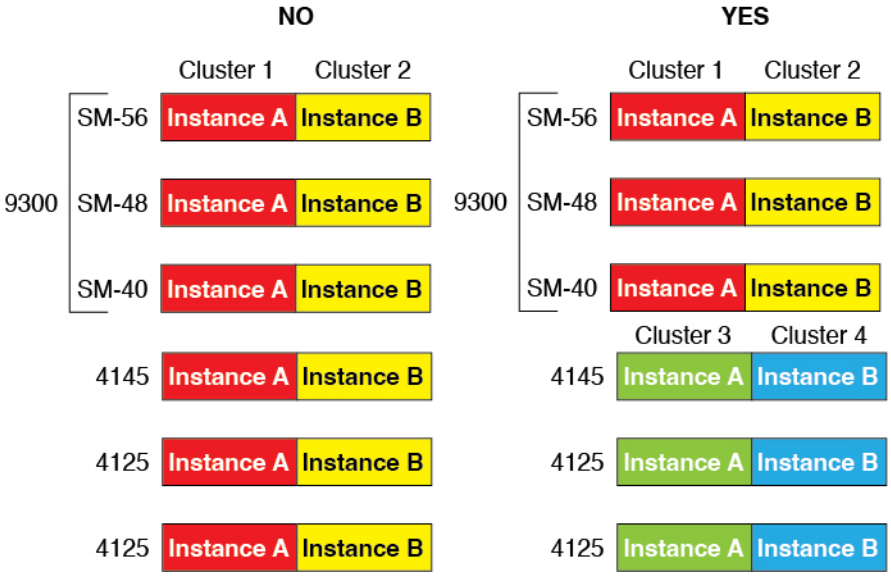

- Maximum 6 nodes—You can use up to six container instances in a cluster.

Switch Requirements

- Be sure to complete the switch configuration and successfully connect all the EtherChannels from the chassis to the switch(es) before you configure clustering on the Firepower 4100/9300 chassis.
- For supported switch characteristics, see Cisco FXOS Compatibility.

Guidelines and Limitations for Logical Devices

See the following sections for guidelines and limitations.

Guidelines and Limitations for Interfaces

VLAN Subinterfaces

- This document discusses *FXOS* VLAN subinterfaces only. You can separately create subinterfaces within the Firewall Threat Defense application. See FXOS Interfaces vs. Application Interfaces, on page 300 for more information.
- Subinterfaces (and the parent interfaces) can only be assigned to container instances.

Note If you assign a parent interface to a container instance, it only passes untagged (non-VLAN) traffic. Do not assign the parent interface unless you intend to pass untagged traffic. For Cluster type interfaces, the parent interface cannot be used.

- Subinterfaces are supported on Data or Data-sharing type interfaces, as well as Cluster type interfaces. If you add subinterfaces to a Cluster interface, you cannot use that interface for a native cluster.
- For multi-instance clustering, FXOS subinterfaces are not supported on Data interfaces. However, subinterfaces are supported for the cluster control link, so you can use either a dedicated EtherChannel or a subinterface of an EtherChannel for the cluster control link. Note that *application*-defined subinterfaces are supported for Data interfaces.
- You can create up to 500 VLAN IDs.
- See the following limitations within the logical device application; keep these limitations in mind when planning your interface allocation.
  - You cannot use subinterfaces for an Firewall Threat Defense inline set or as a passive interface.
  - If you use a subinterface for the failover link, then all subinterfaces on that parent, and the parent itself, are restricted for use as failover links. You cannot use some subinterfaces as failover links, and some as regular data interfaces.

Data-sharing Interfaces

- You cannot use a data-sharing interface with a native instance.
- Maximum 14 instances per shared interface. For example, you can allocate Ethernet1/1 to Instance1 through Instance14.

  Maximum 10 shared interfaces per instance. For example, you can allocate Ethernet1/1.1 through Ethernet1/1.10 to Instance1.

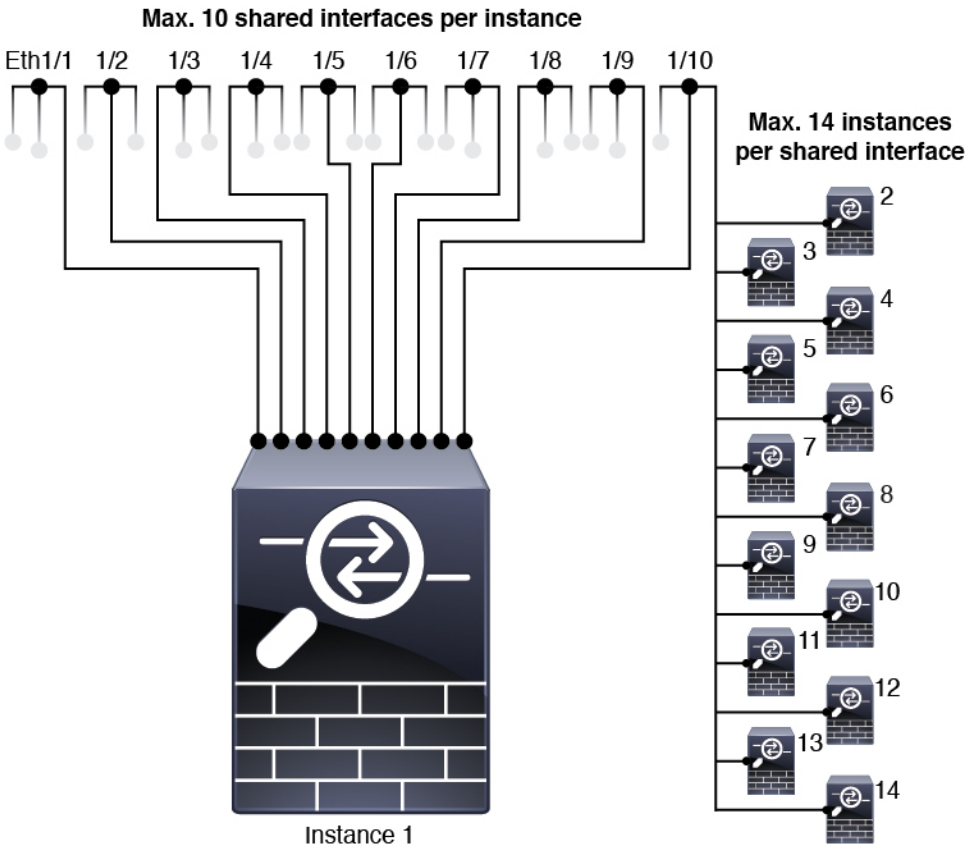

- You cannot use a data-sharing interface in a cluster.
- See the following limitations within the logical device application; keep these limitations in mind when planning your interface allocation.
  - You cannot use a data-sharing interface with a transparent firewall mode device.
  - You cannot use a data-sharing interface with Firewall Threat Defense inline sets or passive interfaces.
  - You cannot use a data-sharing interface for the failover link.

Inline Sets for Firewall Threat Defense

- Supported for physical interfaces (both regular and breakout ports) and EtherChannels. Subinterfaces are not supported.
- Link state propagation is supported.
- Do not enable Hardware Bypass and link state propagation for the same inline set.

Hardware Bypass

- Supported for the Firewall Threat Defense; you can use them as regular interfaces for the ASA.
- The Firewall Threat Defense only supports Hardware Bypass with inline sets.

- Hardware Bypass-capable interfaces cannot be configured for breakout ports.
- You cannot include Hardware Bypass interfaces in an EtherChannel and use them for Hardware Bypass; you can use them as regular interfaces in an EtherChannel.
- Hardware Bypass is not supported with High Availability.
- Do not enable Hardware Bypass and link state propagation for the same inline set.

Default MAC Addresses

For native instances:

Default MAC address assignments depend on the type of interface.

- Physical interfaces—The physical interface uses the burned-in MAC address.
- EtherChannels—For an EtherChannel, all interfaces that are part of the channel group share the same MAC address. This feature makes the EtherChannel transparent to network applications and users, because they only see the one logical connection; they have no knowledge of the individual links. The port-channel interface uses a unique MAC address from a pool; interface membership does not affect the MAC address.

For container instances:

- MAC addresses for all interfaces are taken from a MAC address pool. For subinterfaces, if you decide to manually configure MAC addresses, make sure you use unique MAC addresses for all subinterfaces on the same parent interface to ensure proper classification. See Automatic MAC Addresses for Container Instance Interfaces, on page 319.

General Guidelines and Limitations

Firewall Mode

You can set the firewall mode to routed or transparent in the bootstrap configuration for the Firewall Threat Defense.

High Availability

- Configure high availability within the application configuration.
- You can use any data interfaces as the failover and state links. Data-sharing interfaces are not supported.

Multi-Instance

- Multi-instance capability with container instances is only available for the Firewall Threat Defense using Firewall Management Center.
- For Firewall Threat Defense container instances, a single Firewall Management Center must manage all instances on a security module/engine.
- For Firewall Threat Defense container instances, the following features are not supported:
  - Radware DefensePro link decorator

- Firewall Management Center UCAPL/CC mode
- Flow offload to hardware

Configure Interfaces

By default, physical interfaces are disabled. You can enable interfaces, add EtherChannels, add VLAN subinterfaces, and edit interface properties.

Enable or Disable an Interface

You can change the **Admin State** of each interface to be enabled or disabled. By default, physical interfaces are disabled. For VLAN subinterfaces, the admin state is inherited from the parent interface.

Procedure

Step 1 Choose **Interfaces** to open the Interfaces page.

The Interfaces page shows a visual representation of the currently installed interfaces at the top of the page and provides a listing of the installed interfaces in the table below.

Step 2 To enable the interface, click the disabled **Slider disabled** ( ) so that it changes to the enabled **Slider enabled** ( ).

Click **Yes** to confirm the change. The corresponding interface in the visual representation changes from gray to green.

Step 3 To disable the interface, click the enabled **Slider enabled** ( ) so that it changes to the disabled **Slider disabled** ( ).

Click **Yes** to confirm the change. The corresponding interface in the visual representation changes from green to gray.

Configure a Physical Interface

You can physically enable and disable interfaces, as well as set the interface speed and duplex. To use an interface, it must be physically enabled in FXOS and logically enabled in the application.

Note

- For QSFPH40G-CUxM, auto-negotiation is always enabled by default and you cannot disable it.
- If you replace an SFP on a port with a different SFP module, the speed, duplex, and auto-negotiation of the interface is not updated automatically. You must manually re-configure the interface.

Before you begin

- Interfaces that are already a member of an EtherChannel cannot be modified individually. Be sure to configure settings before you add it to the EtherChannel.

Procedure

Step 1 Choose **Interfaces** to open the Interfaces page.

The **All Interfaces** page shows a visual representation of the currently installed interfaces at the top of the page and provides a listing of the installed interfaces in the table below.

Step 2 Click **Edit** in the row for the interface you want to edit to open the **Edit Interface** dialog box.

Step 3 To enable the interface, check the **Enable** check box. To disable the interface, uncheck the **Enable** check box.

Step 4 Choose the interface **Type**:

See Interface Types, on page 298 for details about interface type usage.

- **Data**
- **Data-sharing**—For container instances only.
- **Mgmt**
- **Firepower-eventing**—For Firewall Threat Defense only.
- **Cluster**—Do not choose the **Cluster** type; by default, the cluster control link is automatically created on Port-channel 48.

Step 5 (Optional) Choose the speed of the interface from the **Speed** drop-down list.

Step 6 (Optional) If your interface supports **Auto Negotiation**, click the **Yes** or **No** radio button.

Step 7 (Optional) Choose the duplex of the interface from the **Duplex** drop-down list.

Step 8 (Optional) Explicitly configure **Debounce Time (ms)**. Enter a value between 0-15000 milli-seconds.

Note
Configuring Debounce Time is not supported on 1G interfaces.

Step 9 Click **OK**.

Add an EtherChannel (Port Channel)

An EtherChannel (also known as a port channel) can include up to 16 member interfaces of the same media type and capacity, and must be set to the same speed and duplex. The media type can be either RJ-45 or SFP; SFPs of different types (copper and fiber) can be mixed. You cannot mix interface capacities (for example 1GB and 10GB interfaces) by setting the speed to be lower on the larger-capacity interface. The Link Aggregation Control Protocol (LACP) aggregates interfaces by exchanging the Link Aggregation Control Protocol Data Units (LACPDUs) between two network devices.

You can configure each physical Data or Data-sharing interface in an EtherChannel to be:

- Active—Sends and receives LACP updates. An active EtherChannel can establish connectivity with either an active or a passive EtherChannel. You should use the active mode unless you need to minimize the amount of LACP traffic.
- On—The EtherChannel is always on, and LACP is not used. An "on" EtherChannel can only establish a connection with another "on" EtherChannel.

Note It may take up to three minutes for an EtherChannel to come up to an operational state if you change its mode from On to Active or from Active to On.

Non-data interfaces only support active mode.

LACP coordinates the automatic addition and deletion of links to the EtherChannel without user intervention. It also handles misconfigurations and checks that both ends of member interfaces are connected to the correct channel group. "On" mode cannot use standby interfaces in the channel group when an interface goes down, and the connectivity and configurations are not checked.

When the Firepower 4100/9300 chassis creates an EtherChannel, the EtherChannel stays in a **Suspended** state for Active LACP mode or a **Down** state for On LACP mode until you assign it to a logical device, even if the physical link is up. The EtherChannel will be brought out of this **Suspended** state in the following situations:

- The EtherChannel is added as a data or management interface for a standalone logical device
- The EtherChannel is added as a management interface or cluster control link for a logical device that is part of a cluster
- The EtherChannel is added as a data interface for a logical device that is part of a cluster and at least one unit has joined the cluster

Note that the EtherChannel does not come up until you assign it to a logical device. If the EtherChannel is removed from the logical device or the logical device is deleted, the EtherChannel will revert to a **Suspended** or **Down** state.

Procedure

Step 1 Choose **Interfaces** to open the Interfaces page.

The **All Interfaces** page shows a visual representation of the currently installed interfaces at the top of the page and provides a listing of the installed interfaces in the table below.

Step 2 Click **Add Port Channel** above the interfaces table to open the **Add Port Channel** dialog box.

Step 3 Enter an ID for the port channel in the **Port Channel ID** field. Valid values are between 1 and 47.

Port-channel 48 is reserved for the cluster control link when you deploy a clustered logical device. If you do not want to use Port-channel 48 for the cluster control link, you can delete it and configure a Cluster type EtherChannel with a different ID.You can add multiple Cluster type EtherChannels and add VLAN subinterfaces for use with multi-instance clustering. For intra-chassis clustering, do not assign any interfaces to the Cluster EtherChannel.

Step 4 To enable the port channel, check the **Enable** check box. To disable the port channel, uncheck the **Enable** check box.

Step 5 Choose the interface **Type**:

See Interface Types, on page 298 for details about interface type usage.

- **Data**
- **Data-sharing**—For container instances only.
- **Mgmt**
- **Firepower-eventing**—For Firewall Threat Defense only.
- **Cluster**

Step 6 Set the required **Admin Speed** for the member interfaces from the drop-down list.

If you add a member interface that is not at the specified speed, it will not successfully join the port channel.

Step 7 For Data or Data-sharing interfaces, choose the LACP port-channel **Mode**, **Active** or **On**.

For non-Data or non-Data-sharing interfaces, the mode is always active.

Step 8 Set the required **Admin Duplex** for the member interfaces, **Full Duplex** or **Half Duplex**.

If you add a member interface that is configured with the specified duplex, it will not successfully join the port channel.

Step 9 To add an interface to the port channel, select the interface in the **Available Interface** list and click **Add Interface** to move the interface to the Member ID list.

You can add up to 16 member interfaces of the same media type and capacity. The member interfaces must be set to the same speed and duplex, and must match the speed and duplex that you configured for this port channel. The media type can be either RJ-45 or SFP; SFPs of different types (copper and fiber) can be mixed. You cannot mix interface capacities (for example 1GB and 10GB interfaces) by setting the speed to be lower on the larger-capacity interface.

Tip
You can add multiple interfaces at one time. To select multiple individual interfaces, click on the desired interfaces while holding down the **Ctrl** key. To select a range of interfaces, select the first interface in the range, and then, while holding down the **Shift** key, click to select the last interface in the range.

Step 10 To remove an interface from the port channel, click the **Delete** button to the right of the interface in the Member ID list.

Step 11 Click **OK**.

Add a VLAN Subinterface for Container Instances

You can add between 250 and 500 VLAN subinterfaces to the chassis, depending on your network deployment. You can add up to 500 subinterfaces to your chassis.

For multi-instance clustering, you can only add subinterfaces to the Cluster-type interface; subinterfaces on data interfaces are not supported.

VLAN IDs per interface must be unique, and within a container instance, VLAN IDs must be unique across all assigned interfaces. You can reuse VLAN IDs on *separate* interfaces as long as they are assigned to different container instances. However, each subinterface still counts towards the limit even though it uses the same ID.

This document discusses *FXOS* VLAN subinterfaces only. You can separately create subinterfaces within the Firewall Threat Defense application. For more information on when to use FXOS subinterfaces vs. application subinterfaces, see FXOS Interfaces vs. Application Interfaces, on page 300.

Procedure

Step 1 Choose **Interfaces** to open the **All Interfaces** tab.

The **All Interfaces** tab shows a visual representation of the currently installed interfaces at the top of the page and provides a listing of the installed interfaces in the table below.

Step 2 Click **Add New** > **Subinterface** to open the **Add Subinterface** dialog box.

Step 3 Choose the interface **Type**:

See Interface Types, on page 298 for details about interface type usage.

- **Data**
- **Data-sharing**
- **Cluster**—If you add subinterfaces to a Cluster interface, you cannot use that interface for a native cluster.

For Data and Data-sharing interfaces: The type is independent of the parent interface type; you can have a Data-sharing parent and a Data subinterface, for example.

Step 4 Choose the parent **Interface** from the drop-down list.

You cannot add a subinterface to a physical interface that is currently allocated to a logical device. If other subinterfaces of the parent are allocated, you can add a new subinterface as long as the parent interface itself is not allocated.

Step 5 Enter a **Subinterface ID**, between 1 and 4294967295.

This ID will be appended to the parent interface ID as *interface_id.subinterface_id*. For example, if you add a subinterface to Ethernet1/1 with the ID of 100, then the subinterface ID will be: Ethernet1/1.100. This ID is not the same as the VLAN ID, although you can set them to match for convenience.

Step 6 Set the **VLAN ID** between 1 and 4095.

Step 7 Click **OK**.

Expand the parent interface to view all subinterfaces under it.

Configure Logical Devices

Add a standalone logical device or a High Availability pair on the Firepower 4100/9300.

For clustering, see Clustering for the Firepower 4100/9300, on page 699.

Add a Resource Profile for Container Instances

To specify resource usage per container instance, create one or more resource profiles. When you deploy the logical device/application instance, you specify the resource profile that you want to use. The resource profile sets the number of CPU cores; RAM is dynamically allocated according to the number of cores, and disk space is set to 40 GB per instance.

- The minimum number of cores is 6.

Note Instances with a smaller number of cores might experience relatively higher CPU utilization than those with larger numbers of cores. Instances with a smaller number of cores are more sensitive to traffic load changes. If you experience traffic drops, try assigning more cores.

- You can assign cores as an even number (6, 8, 10, 12, 14 etc.) up to the maximum.
- The maximum number of cores available depends on the security module/chassis model; see Requirements and Prerequisites for Container Instances, on page 323.

The chassis includes a default resource profile called "Default-Small," which includes the minimum number of cores. You can change the definition of this profile, and even delete it if it is not in use. Note that this profile is created when the chassis reloads and no other profile exists on the system.

Changing the resource profile after you assign it is disruptive. See the following guidelines:

- You cannot change the resource profile settings if it is currently in use. You must disable any instances that use it, then change the resource profile, and finally reenable the instance.
- If you change the resource profile settings after you add the Firewall Threat Defense instance to the Firewall Management Center, then update the inventory for each unit on the Firewall Management Center. Choose **Devices** > **Device Management**, click **Edit** ( ) for the instance, then click **Refresh** () on the **Device** > **Inventory Details** area.
- If you assign a different profile to an instance, it reboots.
- If you assign a different profile to instances in an established high-availability pair, which requires the profile to be the same on both units, you must:
  1. Break high availability.
  2. Assign the new profile to both units.
  3. Re-establish high availability.
- If you assign a different profile to instances in an established cluster, which allows mismatched profiles, then apply the new profile on the data nodes first; after they all come back up, you can apply the new profile to the control node.

Procedure

Step 1 Choose **Platform Settings** > **Resource Profiles** , and click **Add**.

The **Add Resource Profile** dialog box appears.

Step 2 Set the following paramters.

- **Name**—Sets the name of the profile between 1 and 64 characters. Note that you cannot change the name of this profile after you add it.
- **Description**—Sets the description of the profile up to 510 characters.
- **Number of Cores**—Sets the number of cores for the profile, between 6 and the maximum, depending on your chassis, as an even number.

Step 3 Click **OK**.

Add a Standalone Firewall Threat Defense

Standalone logical devices work either alone or in a High Availability pair. On the Firepower 9300 with multiple security modules, you can deploy either a cluster or standalone devices. The cluster must use all modules, so you cannot mix and match a 2-module cluster plus a single standalone device, for example.

You can use native instances on some modules, and container instances on the other module(s).

Before you begin

- Download the application image you want to use for the logical device from Cisco.com, and then upload that image to the Firepower 4100/9300 chassis.

Note For the Firepower 9300, you can install different application types (ASA and Firewall Threat Defense) on separate modules in the chassis. You can also run different versions of an application instance type on separate modules.

- Configure a management interface to use with the logical device. The management interface is required. Note that this management interface is not the same as the chassis management port that is used only for chassis management (and that appears at the top of the **Interfaces** tab as **MGMT**).
- You can later enable management from a data interface; but you must assign a Management interface to the logical device even if you don't intend to use it after you enable data management. See the **configure network management-data-interface** command in the FTD command reference for more information.
- You must also configure at least one Data type interface. Optionally, you can also create a firepower-eventing interface to carry all event traffic (such as web events). See Interface Types, on page 298 for more information.

- For container instances, if you do not want to use the default profile, add a resource profile according to Add a Resource Profile for Container Instances, on page 337.
- For container instances, before you can install a container instance for the first time, you must reinitialize the security module/engine so that the disk has the correct formatting. Choose **Security Modules** or **Security Engine**, and click the **Reinitialize icon**. An existing logical device will be deleted and then reinstalled as a new device, losing any local application configuration. If you are replacing a native instance with container instances, you will need to delete the native instance in any case. You cannot automatically migrate a native instance to a container instance.
- Gather the following information:
  - Interface IDs for this device
  - Management interface IP address and network mask
  - Gateway IP address
  - Firewall Management Center IP address and/or NAT ID of your choosing
  - DNS server IP address
  - Firewall Threat Defense hostname and domain name

Procedure

Step 1 Choose **Logical Devices**.

Step 2 Click **Add** > **Standalone**, and set the following parameters:

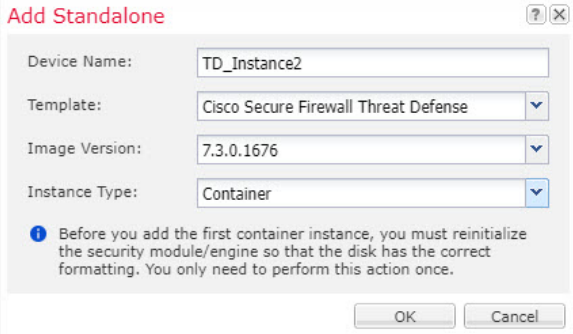

a) Provide a **Device Name**.

This name is used by the chassis supervisor to configure management settings and to assign interfaces; it is not the device name used in the application configuration.

Note
You cannot change this name after you add the logical device.

b) For the **Template**, choose **Cisco Firepower Threat Defense**.
c) Choose the **Image Version**.
d) Choose the **Instance Type**: **Container** or **Native**.

A native instance uses all of the resources (CPU, RAM, and disk space) of the security module/engine, so you can only install one native instance. A container instance uses a subset of resources of the security module/engine, so you can install multiple container instances.

e) Click **OK**.

You see the Provisioning - *device name* window.

Step 3 Expand the **Data Ports** area, and click each interface that you want to assign to the device.

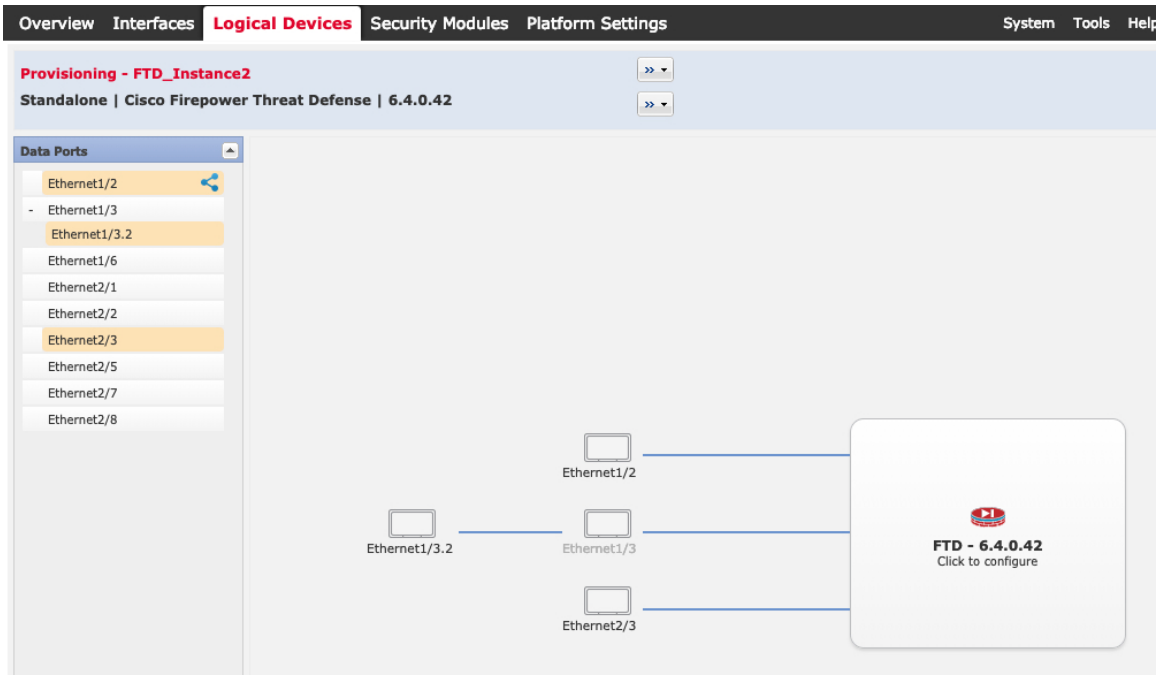

You can only assign data and data-sharing interfaces that you previously enabled on the **Interfaces** page. You will later enable and configure these interfaces in Firewall Management Center, including setting the IP addresses.

You can only assign up to 10 data-sharing interfaces to a container instance. Also, each data-sharing interface can be assigned to at most 14 container instances. A data-sharing interface is indicated by the sharing icon ( ).

Hardware Bypass-capable ports are shown with the following icon: . For certain interface modules, you can enable the Hardware Bypass feature for Inline Set interfaces only (see the Firewall Management Center configuration guide). Hardware Bypass ensures that traffic continues to flow between an inline interface pair during a power outage. This feature can be used to maintain network connectivity in the case of software or hardware failures. If you do not assign both interfaces in a Hardware Bypass pair, you see a warning message to make sure your assignment is intentional. You do not need to use the Hardware Bypass feature, so you can assign single interfaces if you prefer.

Step 4 Click the device icon in the center of the screen.

A dialog box appears where you can configure initial bootstrap settings. These settings are meant for initial deployment only, or for disaster recovery. For normal operation, you can later change most values in the application CLI configuration.

Step 5 On the **General Information** page, complete the following:

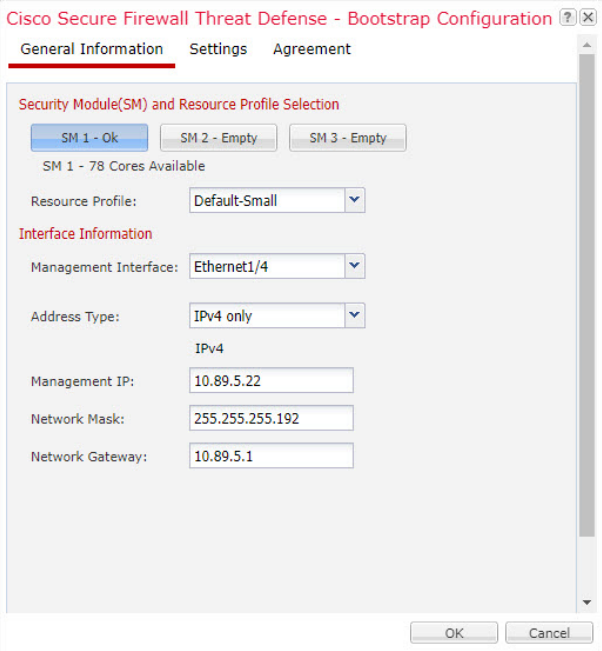

a) (For the Firepower 9300) Under **Security Module Selection** click the security module that you want to use for this logical device.

b) For a container instance, specify the **Resource Profile**.

If you later assign a different resource profile, then the instance will reload, which can take approximately 5 minutes.

Note
If you later assign a different profile to instances in an established high-availability pair, which requires the profile to be the same on both units, you must:

1. Break high availability.
2. Assign the new profile to both units.
3. Re-establish high availability.

c) Choose the **Management Interface**.

This interface is used to manage the logical device. This interface is separate from the chassis management port.

d) Choose the management interface **Address Type**: **IPv4 only**, **IPv6 only**, or **IPv4 and IPv6**.

e) Configure the **Management IP** address.

Set a unique IP address for this interface.

f) Enter a **Network Mask** or **Prefix Length**.

g) Enter a **Network Gateway** address.

Step 6 On the **Settings** tab, complete the following:

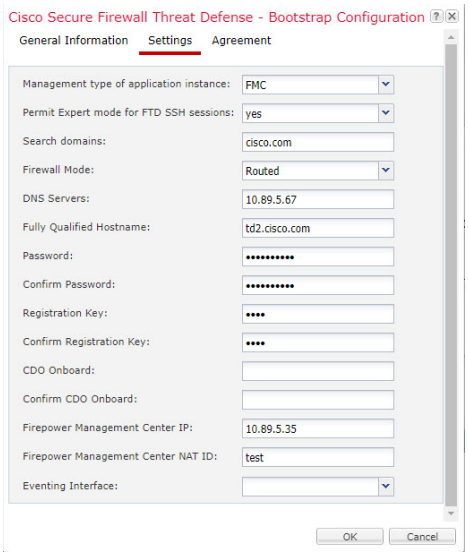

a) For a native instance, in the **Management type of application instance** drop-down list, choose **FMC**.

   Native instances also support Firewall Device Manager as a manager. After you deploy the logical device, you cannot change the manager type.

b) Enter the **Firepower Management Center IP** of the managing Firewall Management Center. If you do not know the Firewall Management Center IP address, leave this field blank and enter a passphrase in the **Firepower Management Center NAT ID** field.

c) For a container instance, **Permit Expert mode from FTD SSH sessions**: **Yes** or **No**. Expert Mode provides Firewall Threat Defense shell access for advanced troubleshooting.

   If you choose **Yes** for this option, then users who access the container instance directly from an SSH sesssion can enter Expert Mode. If you choose **No**, then only users who access the container instance from the FXOS CLI can enter Expert Mode. We recommend choosing **No** to increase isolation between instances.

   Use Expert Mode only if a documented procedure tells you it is required, or if the Cisco Technical Assistance Center asks you to use it. To enter this mode, use the **expert** command in the Firewall Threat Defense CLI.

d) Enter the **Search Domains** as a comma-separated list.

e) Choose the **Firewall Mode**: **Transparent** or **Routed**.

   In routed mode, the Firewall Threat Defense is considered to be a router hop in the network. Each interface that you want to route between is on a different subnet. A transparent firewall, on the other hand, is a Layer 2 firewall that acts like a “bump in the wire,” or a “stealth firewall,” and is not seen as a router hop to connected devices.

   The firewall mode is only set at initial deployment. If you re-apply the bootstrap settings, this setting is not used.

f) Enter the **DNS Servers** as a comma-separated list.

   The Firewall Threat Defense uses DNS if you specify a hostname for the Firewall Management Center, for example.

g) Enter the **Fully Qualified Hostname** for the Firewall Threat Defense.

h) Enter a **Registration Key** to be shared between the Firewall Management Center and the device during registration.

You can choose any text string for this key between 1 and 37 characters; you will enter the same key on the Firewall Management Center when you add the Firewall Threat Defense.

i) Enter a **Password** for the Firewall Threat Defense admin user for CLI access.

j) Choose the **Eventing Interface** on which events should be sent. If not specified, the management interface will be used.

This interface must be defined as a Firepower-eventing interface.

k) For a container instance, set the **Hardware Crypto** as **Enabled** or **Disabled**.

This setting enables TLS crypto acceleration in hardware, and improves performance for certain types of traffic. This feature is enabled by default. You can enable TLS crypto acceleration for up to 16 instances per security module. This feature is always enabled for native instances. To view the percentage of hardware crypto resources allocated to this instance, enter the **show hw-crypto** command.

Step 7 On the **Agreement** tab, read and accept the end user license agreement (EULA).

Step 8 Click **OK** to close the configuration dialog box.

Step 9 Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for the status of the new logical device. When the logical device shows its **Status** as **online**, you can start configuring the security policy in the application.

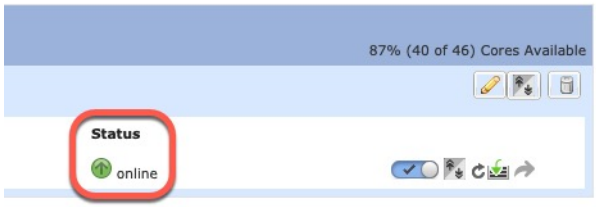

Step 10 See the Firewall Management Center configuration guide to add the Firewall Threat Defense as a managed device and start configuring your security policy.

Add a High Availability Pair

Firewall Threat Defense High Availability (also known as failover) is configured within the application, not in FXOS. However, to prepare your chassis for high availability, see the following steps.

Before you begin

See Requirements and Prerequisites for High Availability, on page 324.

Procedure

Step 1 Allocate the same interfaces to each logical device.

Step 2 Allocate 1 or 2 data interfaces for the failover and state link(s).

These interfaces exchange high availability traffic between the 2 chassis. We recommend that you use a 10 GB data interface for a combined failover and state link. If you have available interfaces, you can use separate failover and state links; the state link requires the most bandwidth. You cannot use the management-type interface for the failover or state link. We recommend that you use a switch between the chassis, with no other device on the same network segment as the failover interfaces.

For container instances, data-sharing interfaces are not supported for the failover link. We recommend that you create subinterfaces on a parent interface or EtherChannel, and assign a subinterface for each instance to use as a failover link. Note that you must use all subinterfaces on the same parent as failover links. You cannot use one subinterface as a failover link and then use other subinterfaces (or the parent interface) as regular data interfaces.

Step 3 Enable High Availability on the logical devices. See High Availability, on page 423.

Step 4 If you need to make interface changes after you enable High Availability, perform the changes on the standby unit first, and then perform the changes on the active unit.

Change an Interface on a Firewall Threat Defense Logical Device

You can allocate or unallocate an interface, or replace a management interface on the Firewall Threat Defense logical device. You can then sync the interface configuration in the Firewall Management Centerthe .

Adding a new interface, or deleting an unused interface has minimal impact on the Firewall Threat Defense configuration. However, deleting an interface that is used in your security policy will impact the configuration. Interfaces can be referenced directly in many places in the Firewall Threat Defense configuration, including access rules, NAT, SSL, identity rules, VPN, DHCP server, and so on. Policies that refer to security zones are not affected. You can also edit the membership of an allocated EtherChannel without affecting the logical device or requiring a sync on the Firewall Management Centerthe .

Deleting an interface will delete any configuration associated with that interface.

Before you begin

- Configure your interfaces, and add any EtherChannels according to Configure a Physical Interface, on page 332 and Add an EtherChannel (Port Channel), on page 333.
- If you want to add an already-allocated interface to an EtherChannel (for example, all interfaces are allocated by default to a cluster), you need to unallocate the interface from the logical device first, then add the interface to the EtherChannel. For a new EtherChannel, you can then allocate the EtherChannel to the device.
- If you want to replace the management or eventing interface with a management EtherChannel, then you need to create the EtherChannel with at least 1 unallocated data member interface, and then replace the current management interface with the EtherChannel. After the Firewall Threat Defense device reboots (management interface changes cause a reboot), and you sync the configuration in the Firewall

Management Centerthe , you can add the (now unallocated) management interface to the EtherChannel as well.

- For clustering or High Availability, make sure you add or remove the interface on all units before you sync the configuration in the Firewall Management Centerthe . We recommend that you make the interface changes on the data/standby unit(s) first, and then on the control/active unit. Note that new interfaces are added in an administratively down state, so they do not affect interface monitoring.
- In mult-instance mode, for changing a sub-interface with an another sub-interface with the same vlan tag, you must first remove all the configuration (including nameif config) of the interface and then unalloacte the interface from Firewall Chassis Manager. Once unallocated, add the new interface and then use sync interfaces from the Firewall Management Center.

Procedure

Step 1 In the Firewall Chassis Manager, choose **Logical Devices**.

Step 2 Click the **Edit** icon at the top right to edit the logical device.

Step 3 Allocate a new data interface by selecting the interface in the **Data Ports** area.

Do not delete any interfaces yet.

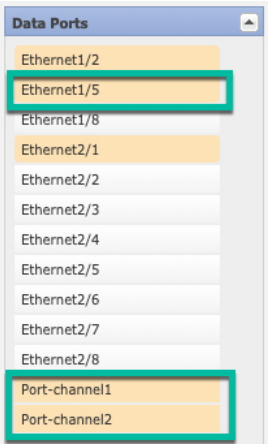

Step 4 Replace the management or eventing interface:

For these types of interfaces, the device reboots after you save your changes.

a) Click the device icon in the center of the page.
b) On the **General** or **Cluster Information** tab, choose the new **Management Interface** from the drop-down list.
c) On the **Settings** tab, choose the new **Eventing Interface** from the drop-down list.
d) Click **OK**.

If you change the IP address of the Management interface, then you must also change the IP address for the device in the Firewall Management Center: go to **Devices** > **Device Management** > **Device/Cluster**. In the **Management** area, set the IP address to match the bootstrap configuration address.

Step 5 Click **Save**.

Step 6 Sync the interfaces in the Firewall Management Center.

a) Log into the Firewall Management Center.
b) Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.
c) Click the **Sync Device** button on the top left of the **Interfaces** page.
d) After the changes are detected, you will see a red banner on the **Interfaces** page indicating that the interface configuration has changed. Click the **Click to know more** link to view the interface changes.
e) If you plan to delete an interface, manually transfer any interface configuration from the old interface to the new interface.

   Because you have not yet deleted any interfaces, you can refer to the existing configuration. You will have additional opportunity to fix the configuration after you delete the old interface and re-run the validation. The validation will show you all locations in which the old interface is still used.

f) Click **Validate Changes** to make sure your policy will still work with the interface changes.

   If there are any errors, you need to change your policy and rerun the validation.

g) Click **Save**.
h) Click **Deploy** > **Deployment**.
i) Select the devices and click **Deploy** to deploy the policy to the assigned devices. The changes are not active until you deploy them.

Step 7 In the Firewall Chassis Manager, unallocate a data interface by de-selecting the interface in the **Data Ports** area.

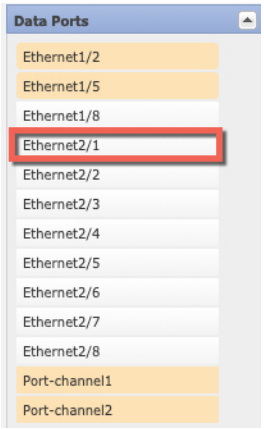

Step 8 Click **Save**.

Step 9 Sync the interfaces again in the Firewall Management Centerthe .

Connect to the Console of the Application

Use the following procedure to connect to the console of the application.

Procedure

Step 1 Connect to the module CLI using a console connection or a Telnet connection.

connect module *slot_number* {**console** | **telnet**}

To connect to the security engine of a device that does not support multiple security modules, always use **1** as the *slot_number*.

The benefits of using a Telnet connection is that you can have multiple sessions to the module at the same time, and the connection speed is faster.

Example:

```
Firepower# connect module 1 console
Telnet escape character is '~'.
Trying 127.5.1.1...
Connected to 127.5.1.1.
Escape character is '~'.

CISCO Serial Over LAN:
Close Network Connection to Exit

Firepower-module1>
```

Step 2 Connect to the application console.

connect ftd *name*

To view the instance names, enter the command without a name.

Example:

```
Firepower-module1> connect ftd ftd1
Connecting to ftd(ftd-native) console... enter exit to return to bootCLI
[...]
>
```

Step 3 Exit the application console to the FXOS module CLI.

- Firewall Threat Defense—Enter **exit**

Step 4 Return to the supervisor level of the FXOS CLI.

Exit the console:

a) Enter **~**

You exit to the Telnet application.

b) To exit the Telnet application, enter:

telnet>**quit**

Exit the Telnet session:

a) Enter **Ctrl-], .**

History for Logical Devices

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Synchronization between the Firewall Threat Defense operational link state and the physical link state	6.7	Any	The chassis can now synchronize the Firewall Threat Defense operational link state with the physical link state for data interfaces. Currently, interfaces will be in an Up state as long as the FXOS admin state is up and the physical link state is up. The Firewall Threat Defense application interface admin state is not considered. Without synchronization from Firewall Threat Defense, data interfaces can be in an Up state physically before the Firewall Threat Defense application has completely come online, for example, or can stay Up for a period of time after you initiate Firewall Threat Defense shutdown. For inline sets, this state mismatch can result in dropped packets because external routers may start sending traffic to the Firewall Threat Defense before the Firewall Threat Defense can handle it. This feature is disabled by default, and can be enabled per logical device in FXOS. **Note** This feature is not supported for clustering, container instances, or Firewall Threat Defense with a Radware vDP decorator. It is also not supported for the ASA. New/Modified Firepower Chassis Manager screens: **Logical Devices** > **Enable Link State** New/Modified FXOS commands: **set link-state-sync enabled**, **show interface expand detail**
Firewall Threat Defense configuration backup and restore using Firewall Management Center for container instances	6.7	Any	You can now use the Firewall Management Center backup/restore tool on Firewall Threat Defense container instances. New/Modified Firewall Management Center screens: **System** > **Tools** > **Backup/Restrore** > **Managed Device Backup** New/Modified Firewall Threat Defense CLI commands: **restore** Supported platforms: Firepower 4100/9300 **Note** Requires FXOS 2.9.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Support for VLAN subinterfaces on a Cluster type interface (multi-instance use only)	6.6	Any	For use with multi-instance clusters, you can now create VLAN subinterfaces on cluster type interfaces. Because each cluster requires a unique cluster control link, VLAN subinterfaces provide a simple method to fulfill this requirement. You can alternatively assign a dedicated EtherChannel per cluster. Multiple cluster interfaces are now allowed. New/Modified Firepower Chassis Manager screens: **Interfaces** > **All Interfaces** > **Add New** drop-down menu**Subinterface** > **Type** field New/Modified FXOS commands: **set port-type cluster** **Note** Requires FXOS 2.8.1.
Firewall Threat Defense on the Firepower 4112	6.6	Any	We introduced the Firepower 4112. **Note** Requires FXOS 2.8.1.
TLS crypto acceleration for multiple container instances	6.5	Any	TLS crypto acceleration is now supported on multiple container instances (up to 16) on a Firepower 4100/9300 chassis. Previously, you could enable TLS crypto acceleration for only *one* container instance per module/security engine. New instances have this feature enabled by default. However, the upgrade does *not* enable acceleration on existing instances. Instead, use the **enter hw-crypto** and then the **set admin-state enabled** FXOS commands. New/Modified Firepower Chassis Manager screens: **Logical Devices** > **Add Device** > **Settings** > **Hardware Crypto** drop-down menu **Note** Requires FXOS 2.7.1.
Firewall Threat Defense on the Firepower 4115, 4125, and 4145	6.4	Any	We introduced the Firepower 4115, 4125, and 4145. **Note** Requires FXOS 2.6.1.157.
Firepower 9300 SM-40, SM-48, and SM-56 support	6.4	Any	We introduced the following three security modules: SM-40, SM-48, and SM-56. **Note** Requires FXOS 2.6.1.157.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Support for ASA and Firewall Threat Defense on separate modules of the same Firepower 9300	6.4	Any	You can now deploy ASA and Firewall Threat Defense logical devices on the same Firepower 9300. **Note** Requires FXOS 2.6.1.157.
Support for SSL hardware acceleration on one Firewall Threat Defense container instance on a module/security engine	6.4	Any	You can now enable SSL hardware acceleration for one container instance on a module/security engine. SSL hardware acceleration is disabled for other container instances, but enabled for native instances. New/Modified FXOS commands: **config hwCrypto enable** No modified screens. **Note** Requires FXOS 2.6.1.157.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Multi-instance capability for Firewall Threat Defense on the Firepower 4100/9300	6.3	Any	You can now deploy multiple logical devices, each with a Firewall Threat Defense container instance, on a single security engine/module. Formerly, you could only deploy a single native application instance. To provide flexible physical interface use, you can create VLAN subinterfaces in FXOS and also share interfaces between multiple instances. Resource management lets you customize performance capabilities for each instance. You can use High Availability using a container instance on 2 separate chassis. Clustering is not supported. **Note** Multi-instance capability is similar to ASA multiple context mode, although the implementation is different. Multiple context mode is not available on the Firewall Threat Defense. New/Modified Firewall Management Center screens: • **Devices** > **Device Management** > **Edit** icon **Interfaces** tab New/Modified Firepower Chassis Manager screens: • **Overview** > **Devices** • **Interfaces** > **All Interfaces** > **Add New** drop-down menu **Subinterface** • **Interfaces** > **All Interfaces** > **Type** • **Logical Devices** > **Add Device** • **Platform Settings** > **Mac Pool** • **Platform Settings** > **Resource Profiles** New/Modified FXOS commands: **connect ftd** *name*, **connect module telnet, create bootstrap-key PERMIT_EXPERT_MODE, create resource-profile, create subinterface, scope auto-macpool, set cpu-core-count, set deploy-type, set port-type data-sharing, set prefix, set resource-profile-name, set vlan, scope app-instance ftd** *name***, show cgroups container, show interface, show mac-address, show subinterface, show tech-support module app-instance, show version** Supported platforms: Firepower 4100/9300

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Cluster control link customizable IP Address for the Firepower 4100/9300	6.3	Any	By default, the cluster control link uses the 127.2.0.0/16 network. You can now set the network when you deploy the cluster in FXOS. The chassis auto-generates the cluster control link interface IP address for each unit based on the chassis ID and slot ID: 127.2.*chassis_id*.*slot_id*. However, some networking deployments do not allow 127.2.0.0/16 traffic to pass. Therefore, you can now set a custom /16 subnet for the cluster control link in FXOS except for loopback (127.0.0.0/8) and multicast (224.0.0.0/4) addresses. New/modified Firepower Chassis Manager screens: • **Logical Devices** > **Add Device** > **Cluster Information** > **CCL Subnet IP** field New/Modified FXOS commands: **set cluster-control-link network** Supported platforms: Firepower 4100/9300
Support for data EtherChannels in On mode	6.3	Any	You can now set data and data-sharing EtherChannels to either Active LACP mode or to On mode. Other types of EtherChannels only support Active mode. New/Modified Firepower Chassis Manager screens: • **Interfaces** > **All Interfaces** > **Edit Port Channel** > **Mode** New/Modified FXOS commands: **set port-channel-mode** Supported platforms: Firepower 4100/9300
Support for EtherChannels in Firewall Threat Defense inline sets	6.2	Any	You can now use EtherChannels in a Firewall Threat Defense inline set. Supported platforms: Firepower 4100/9300
Inter-chassis clustering for 6 Firewall Threat Defense modules	6.2	Any	You can now enable inter-chassis clustering for the Firewall Threat Defense. You can include up to 6 modules in up to 6 chassis. New/modified Firepower Chassis Manager screens: • **Logical Devices** > **Configuration** Supported platforms: Firepower 4100/9300
Hardware bypass support on the Firepower 4100/9300 for supported network modules	6.1	Any	Hardware Bypass ensures that traffic continues to flow between an inline interface pair during a power outage. This feature can be used to maintain network connectivity in the case of software or hardware failures. New/Modified screens: • **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** Supported platforms: Firepower 4100/9300

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Inline set link state propagation support for the Firewall Threat Defense	6.1	Any	When you configure an inline set in the Firewall Threat Defense application and enable link state propagation, the Firewall Threat Defense sends inline set membership to the FXOS chassis. Link state propagation means that the chassis automatically brings down the second interface in the inline interface pair when one of the interfaces in an inline set goes down. New/Modified FXOS commands: **show fault |grep link-down**, **show interface detail** Supported platforms: Firepower 4100/9300
Support for intra-chassis clustering on the Firewall Threat Defense on the Firepower 9300	6.0.1	Any	The Firepower 9300 supports intra-chassis clustering with the Firewall Threat Defense application. New/Modified Firepower Chassis Manager screens: • **Logical Devices** > **Configuration** New/Modified FXOS commands: **enter mgmt-bootstrap ftd, enter bootstrap-key FIREPOWER_MANAGER_IP, enter bootstrap-key FIREWALL_MODE, enter bootstrap-key-secret REGISTRATION_KEY, enter bootstrap-key-secret PASSWORD, enter bootstrap-key FQDN, enter bootstrap-key DNS_SERVERS, enter bootstrap-key SEARCH_DOMAINS, enter ipv4 firepower, enter ipv6 firepower, set value, set gateway, set ip, accept-license-agreement** Supported platforms: Firepower 4100/9300

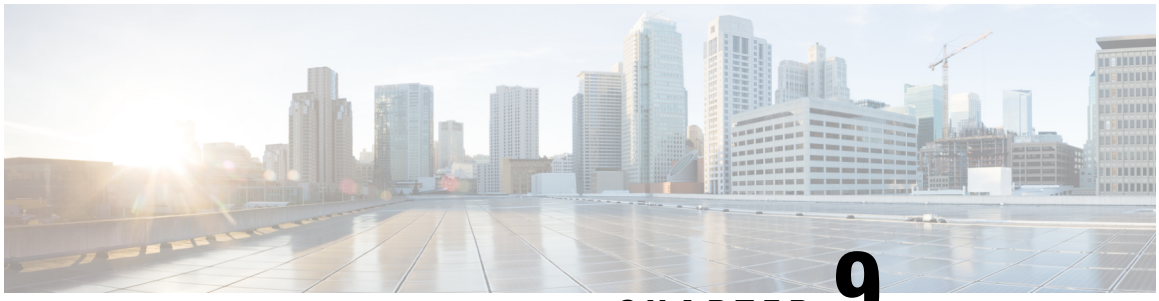

CHAPTER 9

Multi-Instance Mode for the Secure Firewall 3100/4200

You can deploy the Secure Firewall 3100/4200 as a single device (*appliance mode*) or as multiple container instances (*multi-instance mode*). This chapter describes how to deploy the device in multi-instance mode.

- About Multi-Instance Mode, on page 355
- Licenses for Instances, on page 368
- Requirements and Prerequisites for Instances, on page 368
- Guidelines and Limitations for Instances, on page 370
- Configure Instances, on page 372
- Monitoring Multi-Instance Mode, on page 418
- History for Multi-Instance Mode, on page 421

About Multi-Instance Mode

In multi-instance mode, you can deploy multiple container instances on a single chassis that act as completely independent devices.

Multi-Instance Mode vs. Appliance Mode

You can run the device in either multi-instance mode or appliance mode.

Appliance Mode

Appliance mode is the default. The device runs the native Firewall Threat Defense image and acts as a single device. The only chassis-level configuration available (on the **Chassis Manager** page) is for network module management (breakout ports or enabling/disabling a network module).

Multi-Instance Mode

If you change to multi-instance mode, the device runs the Secure Firewall eXtensible Operating System (FXOS) on the chassis, while each instance runs separate Firewall Threat Defense images. You can configure the mode using the FXOS CLI.

Because multiple instances run on the same chassis, you need to perform chassis-level management of:

- CPU and memory resources using resource profiles.

- Interface configuration and assignment.
- Deployment and monitoring of instances.

For a multi-instance device, you add the *chassis* to the Firewall Management Center and configure chassis-level settings on the **Chassis Manager** page.

Chassis Management Interface

Chassis Management

The chassis uses the dedicated Management interface on the device. Multi-instance mode does not support using a data interface for chassis management or DHCP addressing for the Management interface.

You can only configure the chassis Management interface at the Firewall Threat Defense CLI (at initial setup) or FXOS CLI (after you convert to multi-instance mode). See Change Chassis Management Settings at the FXOS CLI, on page 415 to change Management interface settings in multi-instance mode.

Note By default, SSH is not allowed to this interface in multi-instance mode unless you enable the SSH server and an SSH access list. This difference means that you can connect to the application-mode threat defense Management interface using SSH, but after you convert to multi-instance mode, you can no longer connect using SSH by default. See Configure SSH and SSH Access List, on page 397.

Instance Management

All instances share the chassis management interface, and each instance has its own IP address on the Management network. After you add the instance and specify the IP address, you can make changes to the network settings at the Firewall Threat Defense CLI.

The instance Management IP address allows SSH by default.

Instance Eventing Interface

The Secure Firewall 4200 includes a second dedicated interface, Management 1/2, that you can use for events. You can configure this interface at the Firewall Threat Defense CLI in each instance. Assign an IP address on the same network for each instance. See Configure an Event Interface, on page 26.

Instance Interfaces

To provide flexible physical interface use for instances, you can create VLAN subinterfaces on the chassis and also share interfaces (VLAN or physical) between multiple instances. See Shared Interface Scalability, on page 359 and Configure a Subinterface, on page 383.

Note This chapter discusses *chassis* VLAN subinterfaces only. You can separately create subinterfaces within the Firewall Threat Defense instance. See Chassis Interfaces vs. Instance Interfaces, on page 357 for more information.

Interface Types

Physical interfaces, VLAN subinterfaces, and EtherChannel interfaces can be one of the following types:

- Data—Use for regular data or the failover link. Data interfaces cannot be shared between instances, and instances cannot communicate over the backplane to other instances. For traffic on Data interfaces, all traffic must exit the chassis on one interface and return on another interface to reach another instance. You can add VLAN subinterfaces to a data interface to provide separate failover links per High Availability pair.
- Data-sharing—Use for regular data. These data interfaces can be shared by one or more instances. Each instance can communicate over the backplane with all other instances that share this interface. Shared interfaces can affect the number of instances you can deploy. Shared interfaces are not supported for bridge group member interfaces (in transparent mode or routed mode), inline sets, passive interfaces, or failover links.

Chassis Interfaces vs. Instance Interfaces

At the chassis level, you manage the basic Ethernet settings of physical interfaces, VLAN subinterfaces for instances, and EtherChannel interfaces. Within the instance, you configure higher level settings. For example, you can only create EtherChannels in the chassis; but you can assign an IP address to the EtherChannel within the instance.

The following sections describe the interaction between the chassis and the instance for interfaces.

VLAN Subinterfaces

You can create VLAN subinterfaces within the instance, just as you would for any device.

You can *also* create VLAN subinterfaces in the chassis. The instance-defined subinterfaces are not subject to the chassis limit. Choosing in which location to create subinterfaces depends on your network deployment and personal preference. For example, to share a subinterface, you must create the subinterface on the chassis. Another scenario that favors chassis subinterfaces comprises allocating separate subinterface groups on a single interface to multiple instances. For example, you want to use Port-channel1 with VLAN 2–11 on instance A, VLAN 12–21 on instance B, and VLAN 22–31 on instance C. If you create these subinterfaces in the instance, then you would have to share the parent interface in the chassis, which may not be desirable. See the following illustration that shows the three ways you can accomplish this scenario:

Figure 130: VLANs in the Chassis vs. the Instance

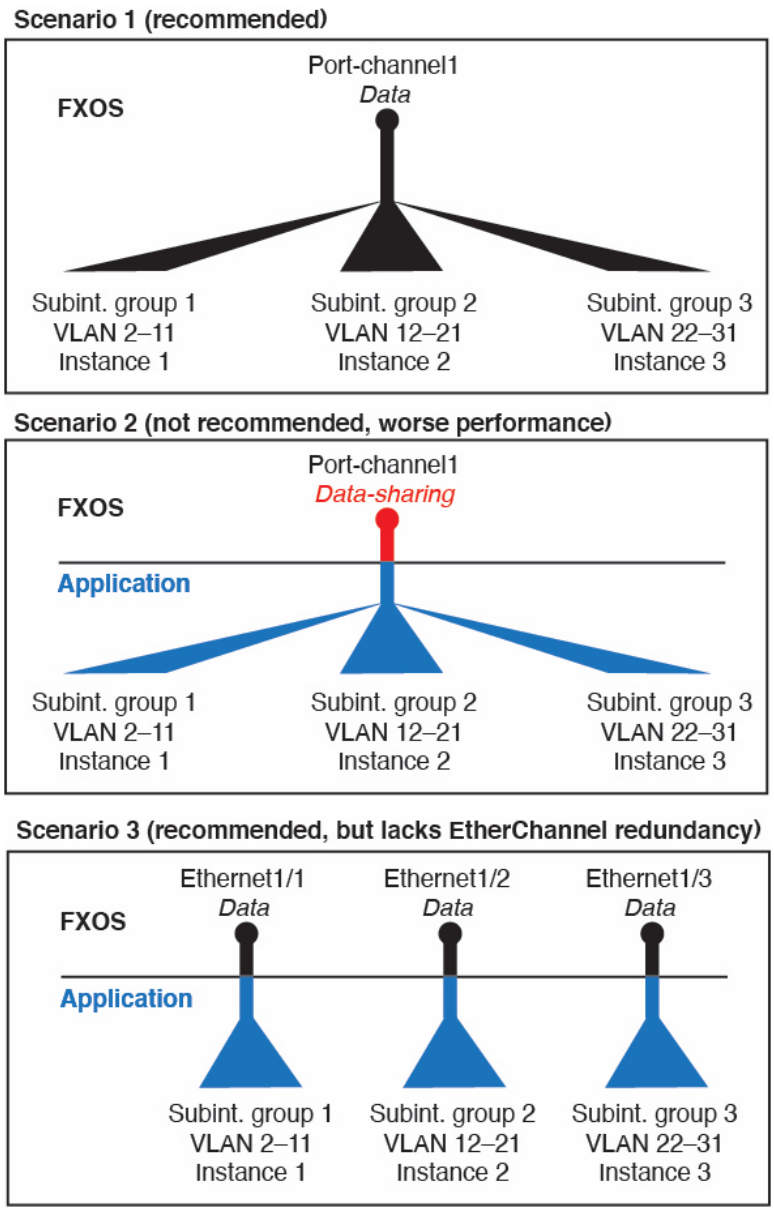

Independent Interface States in the Chassis and in the Instance

You can administratively enable and disable interfaces in both the chassis and in the instance. For an interface to be operational, the interface must be enabled in both locations. Because the interface state is controlled independently, you may have a mismatch between the chassis and instance.

The default state of an interface within the instance depends on the type of interface. For example, the physical interface or EtherChannel is disabled by default within the instance, but a subinterface is enabled by default.

Shared Interface Scalability

Instances can share data-sharing type interfaces. This capability lets you conserve physical interface usage as well as support flexible networking deployments. When you share an interface, the chassis uses unique MAC addresses to forward traffic to the correct instance. However, shared interfaces can cause the forwarding table to grow large due to the need for a full mesh topology within the chassis (every instance must be able to communicate with every other instance that is sharing the same interface). Therefore, there are limits to how many interfaces you can share.

In addition to the forwarding table, the chassis maintains a VLAN group table for VLAN subinterface forwarding. You can create up to 500 VLAN subinterfaces.

See the following limits for shared interface allocation:

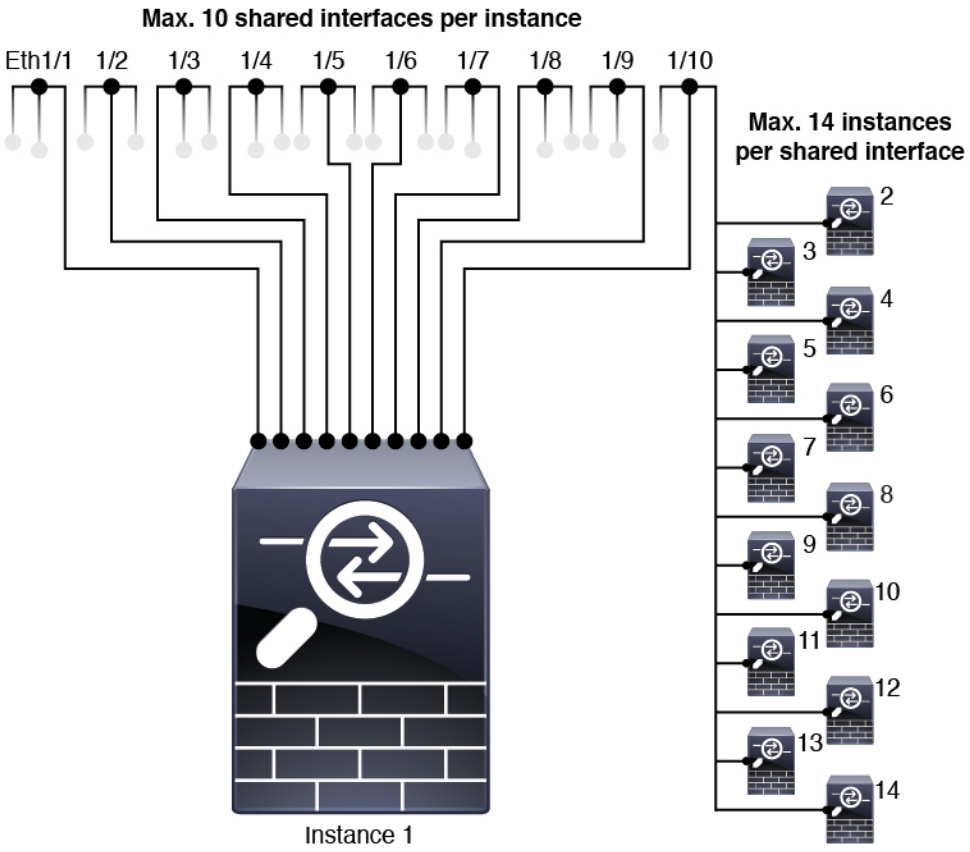

Shared Interface Best Practices

For optimal scalability of the forwarding table, share as few interfaces as possible. Instead, you can create up to 500 VLAN subinterfaces on one or more physical interfaces and then divide the VLANs among the container instances.

When sharing interfaces, follow these practices in the order of most scalable to least scalable:

1. Best—Share subinterfaces under a single parent, and use the same set of subinterfaces with the same group of instances.

   For example, create a large EtherChannel to bundle all of your like-kind interfaces together, and then share subinterfaces of that EtherChannel: Port-Channel1.2, 3, and 4 instead of Port-Channel2,

Port-Channel3, and Port-Channel4. When you share subinterfaces from a single parent, the VLAN group table provides better scaling of the forwarding table than when sharing physical/EtherChannel interfaces or subinterfaces across parents.

Figure 131: Best: Shared Subinterface Group on One Parent

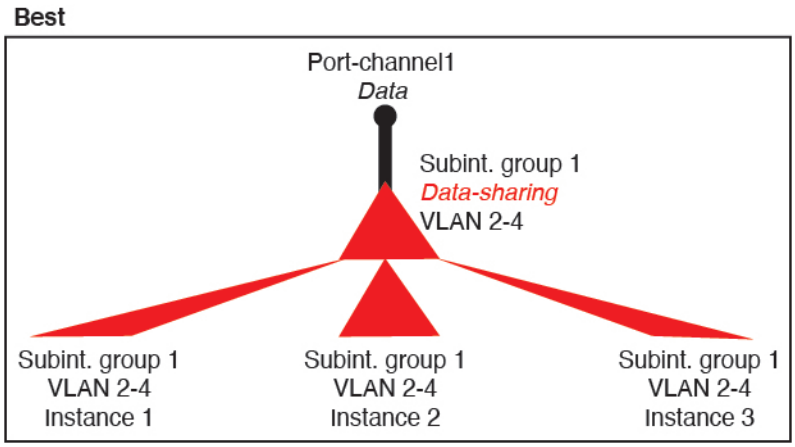

If you do not share the same set of subinterfaces with a group of instances, your configuration can cause more resource usage (more VLAN groups). For example, share Port-Channel1.2, 3, and 4 with instances 1, 2, and 3 (one VLAN group) instead of sharing Port-Channel1.2 and 3 with instances 1 and 2, while sharing Port-Channel1.3 and 4 with instance 3 (two VLAN groups).

Figure 132: Good: Sharing Multiple Subinterface Groups on One Parent

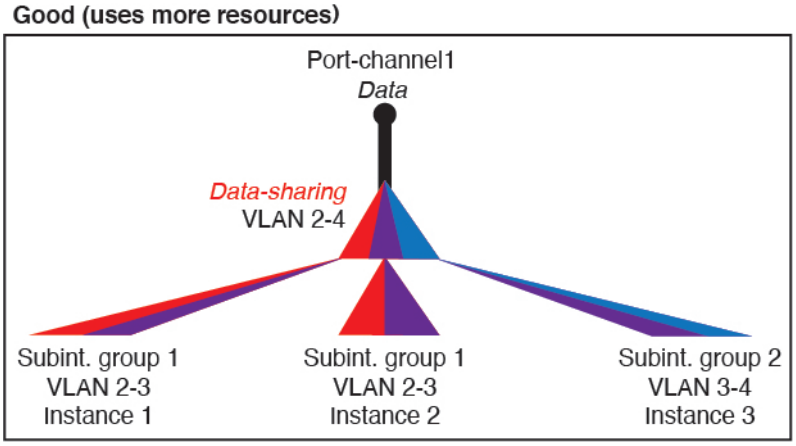

2. Fair—Share subinterfaces across parents.

   For example, share Port-Channel1.2, Port-Channel2.3, and Port-Channel3.4 instead of Port-Channel2, Port-Channel4, and Port-Channel4. Although this usage is not as efficient as only sharing subinterfaces on the same parent, it still takes advantage of VLAN groups.

Figure 133: Fair: Shared Subinterfaces on Separate Parents

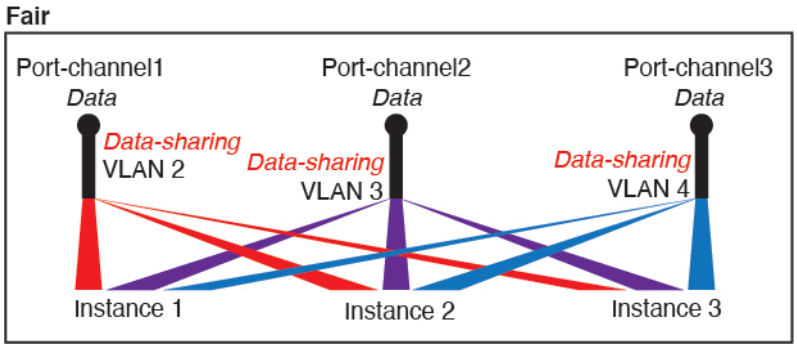

3. Worst—Share individual parent interfaces (physical or EtherChannel).

   This method uses the most forwarding table entries.

Figure 134: Worst: Shared Parent Interfaces

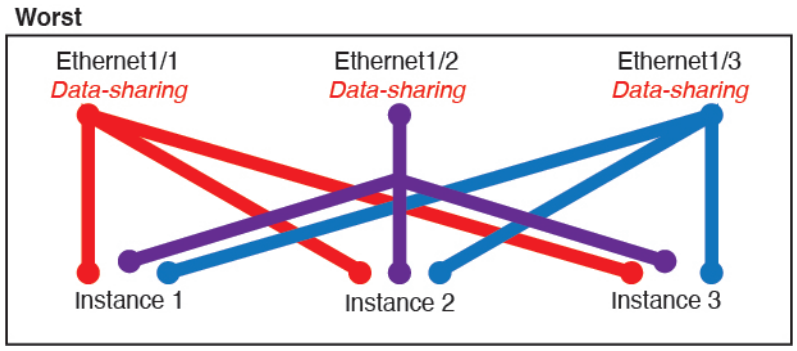

How the Chassis Classifies Packets

Each packet that enters the chassis must be classified, so that the chassis can determine to which instance to send a packet.

- Unique Interfaces—If only one instance is associated with the ingress interface, the chassis classifies the packet into that instance. For bridge group member interfaces (in transparent mode or routed mode), inline sets, or passive interfaces, this method is used to classify packets at all times.
- Unique MAC Addresses—The chassis automatically generates unique MAC addresses for all interfaces, including shared interfaces. If multiple instances share an interface, then the classifier uses unique MAC addresses assigned to the interface in each instance. An upstream router cannot route directly to an instance without unique MAC addresses. You can also set the MAC addresses manually when you configure each interface within the application.

Note If the destination MAC address is a multicast or broadcast MAC address, the packet is duplicated and delivered to each instance.

Classification Examples

Packet Classification with a Shared Interface Using MAC Addresses

The following figure shows multiple instances sharing an outside interface. The classifier assigns the packet to Instance C because Instance C includes the MAC address to which the router sends the packet.

Figure 135: Packet Classification with a Shared Interface Using MAC Addresses

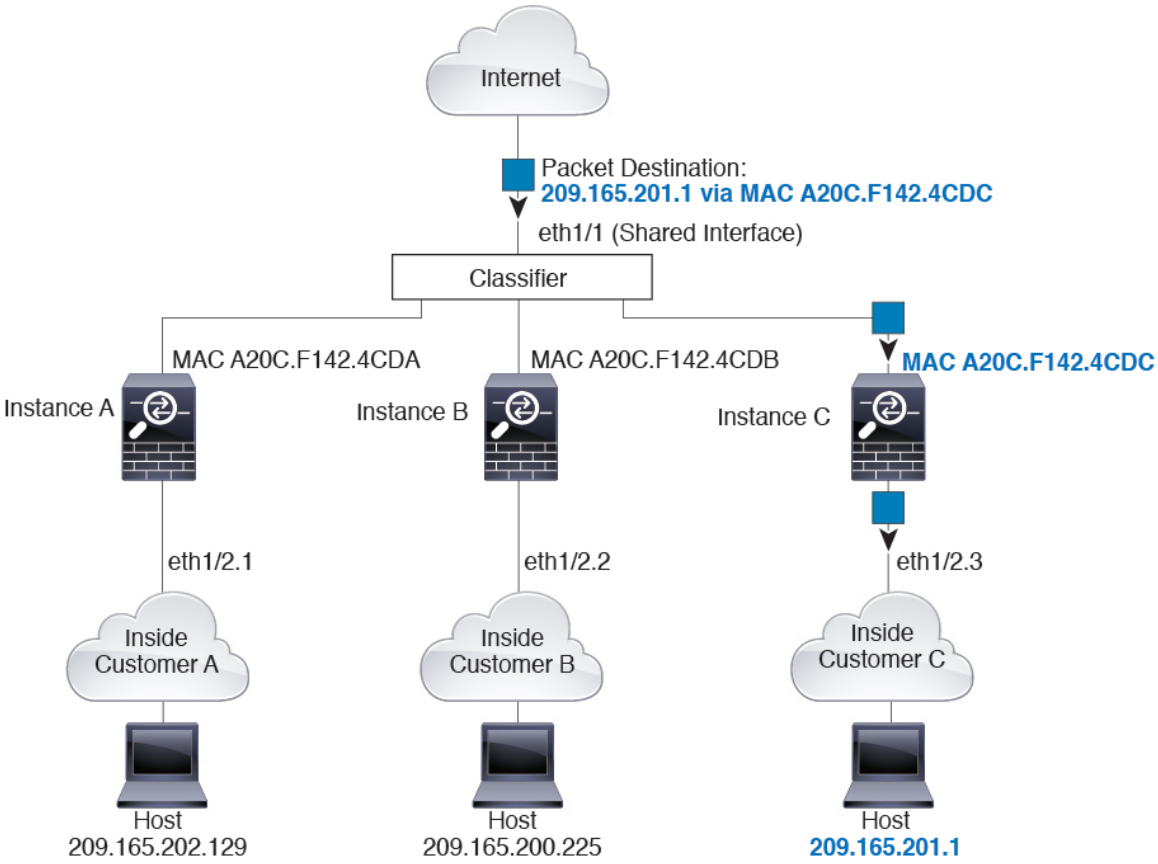

Incoming Traffic from Inside Networks

Note that all new incoming traffic must be classified, even from inside networks. The following figure shows a host on the Instance C inside network accessing the internet. The classifier assigns the packet to Instance C because the ingress interface is Ethernet 1/2.3, which is assigned to Instance C.

Figure 136: Incoming Traffic from Inside Networks

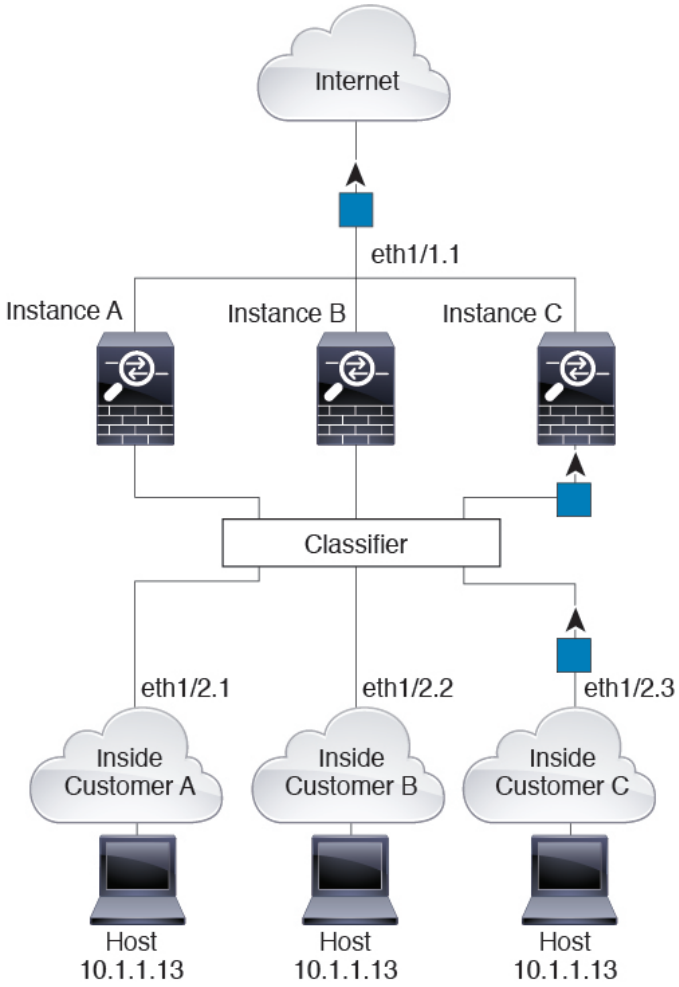

Transparent Firewall Instances

For transparent firewalls, you must use unique interfaces. The following figure shows a packet destined to a host on the Instance C inside network from the internet. The classifier assigns the packet to Instance C because the ingress interface is Ethernet 1/2.3, which is assigned to Instance C.

Figure 137: Transparent Firewall Instances

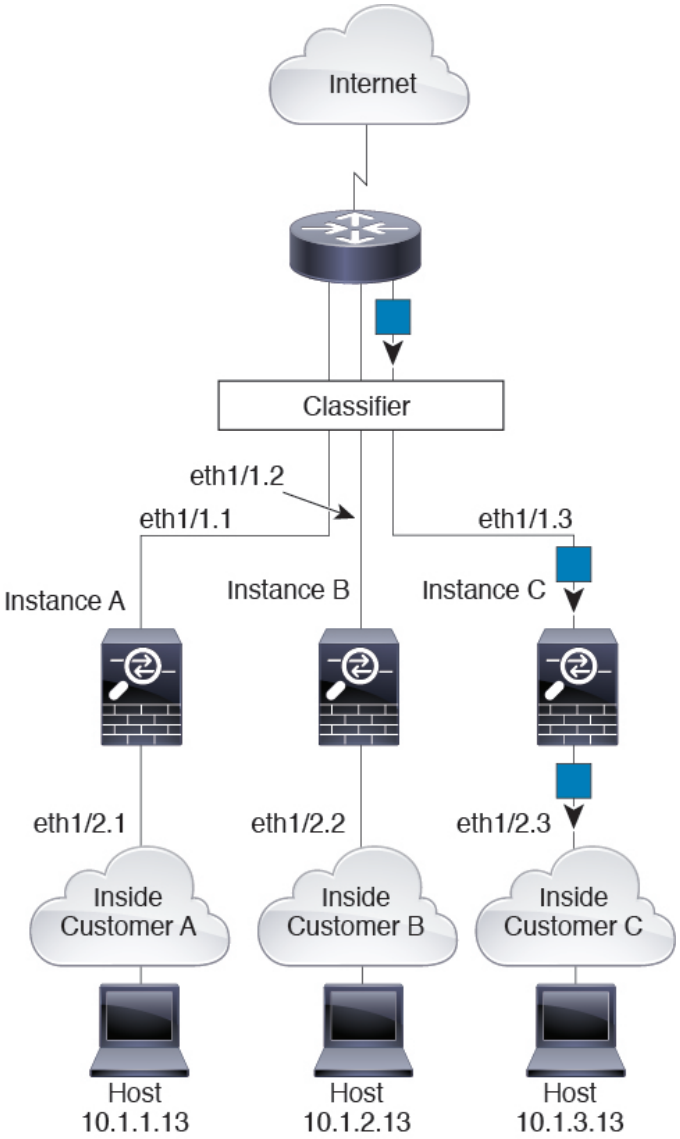

Inline Sets

For inline sets, you must use unique interfaces and they must be physical interfaces or EtherChannels. The following figure shows a packet destined to a host on the Instance C inside network from the internet. The classifier assigns the packet to Instance C because the ingress interface is Ethernet 1/5, which is assigned to Instance C.

Figure 138: Inline Sets

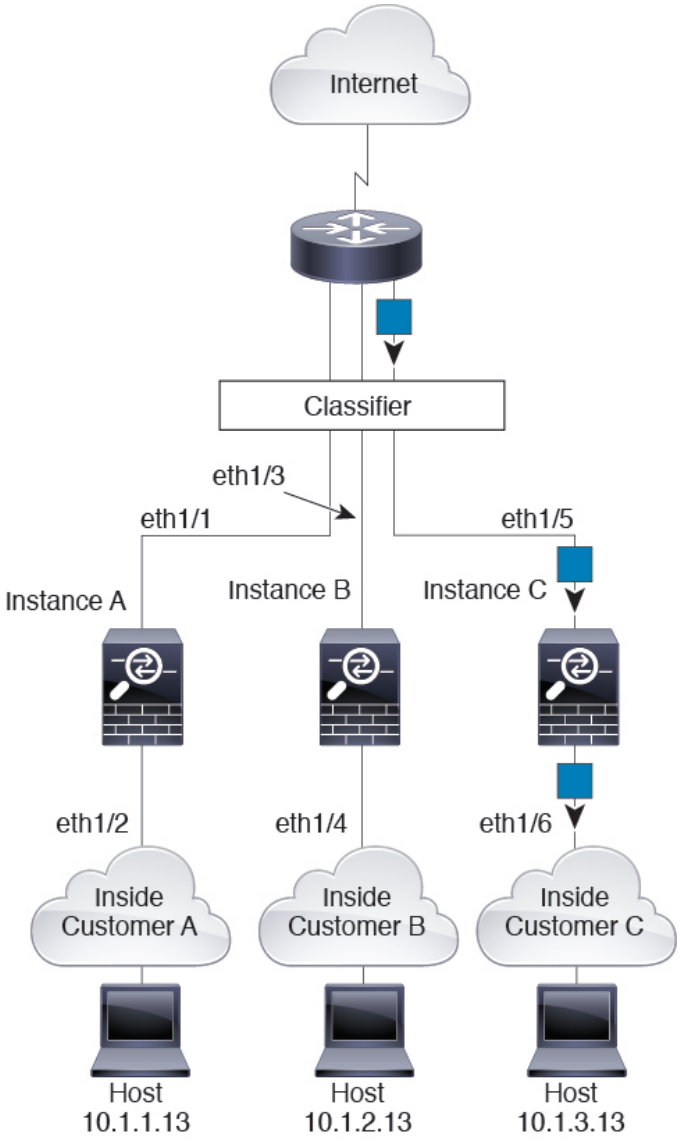

Cascading Instances

Placing an instance directly in front of another instance is called *cascading instances*; the outside interface of one instance is the same interface as the inside interface of another instance. You might want to cascade instances if you want to simplify the configuration of some instances by configuring shared parameters in the top instance.

The following figure shows a gateway instance with two instances behind the gateway.

Figure 139: Cascading Instances

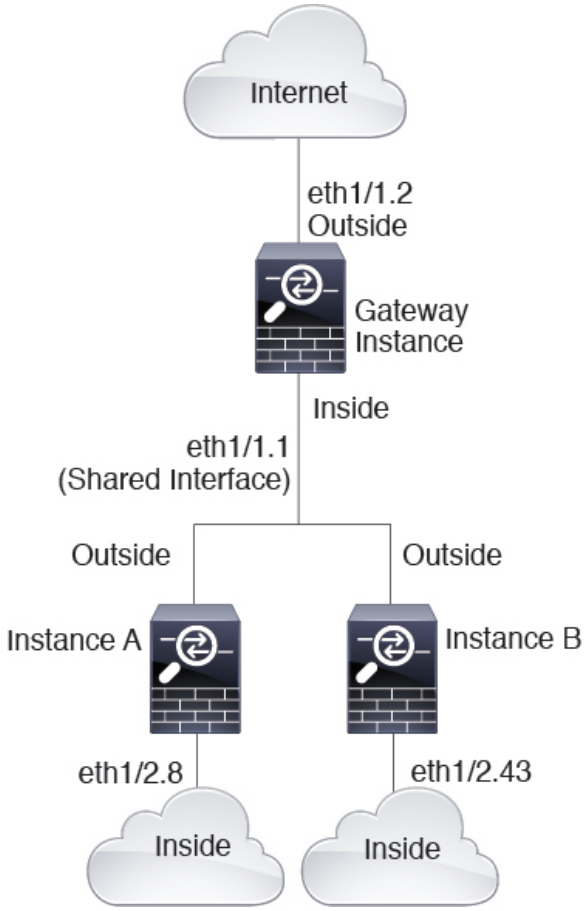

Note Do not use cascading instances (using a shared interface) with High Availability. After a failover occurs and the standby unit rejoins, MAC addresses can overlap temporarily and cause an outage. You should instead use unique interfaces for the gateway instance and inside instance using an external switch to pass traffic between the instances.

Typical Multi-Instance Deployment

The following example includes three container instances in routed firewall mode. They include the following interfaces:

- Management—All instances and the chassis use the dedicated Management interface. Within each instance (and the chassis), the interface uses a unique IP address on the same management network.
- Inside—Each instance uses a subinterface on Port-Channel1 (data type). This EtherChannel includes two 10 Gigibit Ethernet interfaces. Each subinterface is on a separate network.

- Outside—All instances use the Port-Channel2 interface (data-sharing type). This EtherChannel includes two 10 Gigibit Ethernet interfaces. Within each application, the interface uses a unique IP address on the same outside network.
- Failover—Each instance uses a subinterface on Port-Channel3 (data type). This EtherChannel includes two 10 Gigibit Ethernet interfaces. Each subinterface is on a separate network.

Figure 140: Typical Multi-Instance Deployment

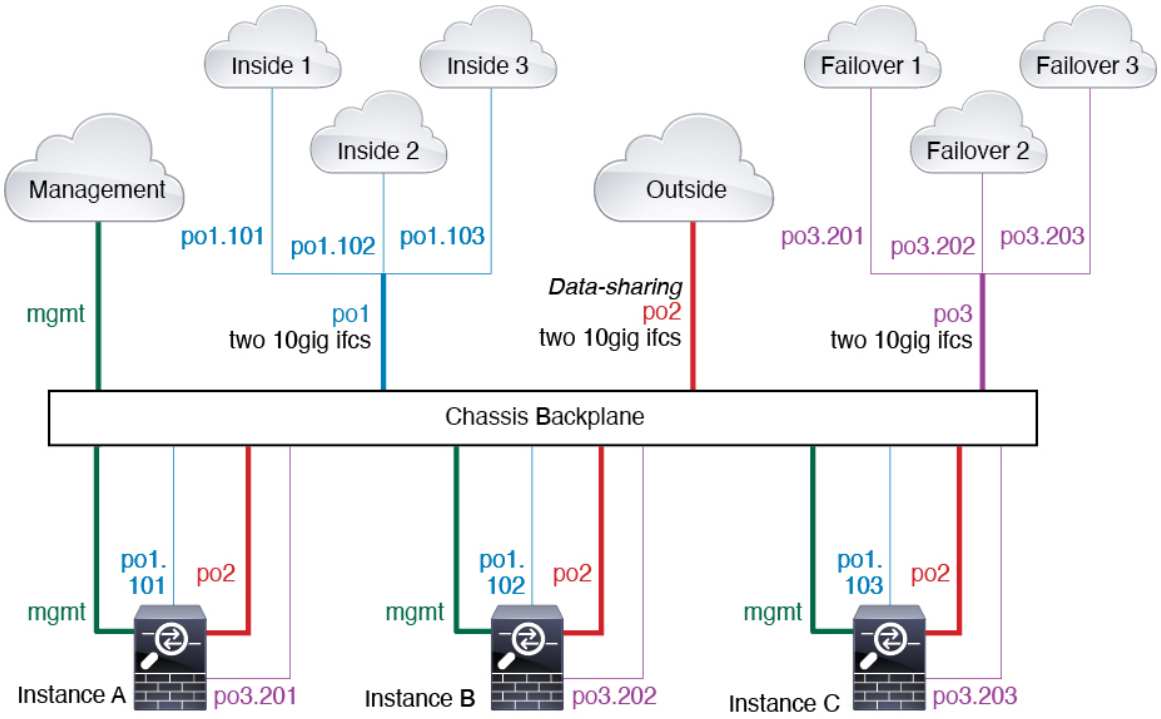

Automatic MAC Addresses for Instance Interfaces

The chassis automatically generates MAC addresses for instance interfaces, and guarantees that a shared interface in each instance uses a unique MAC address.

If you manually assign a MAC address to a shared interface within the instance, then the manually-assigned MAC address is used. If you later remove the manual MAC address, the autogenerated address is used. In the rare circumstance that the generated MAC address conflicts with another private MAC address in your network, we suggest that you manually set the MAC address for the interface within the instance.

Because autogenerated addresses start with A2, you should not start manual MAC addresses with A2 due to the risk of overlapping addresses.

The chassis generates the MAC address using the following format:

A2*xx.yyzz.zzzz*

Where *xx.yy* is a user-defined prefix or a system-defined prefix, and *zz.zzzz* is an internal counter generated by the chassis. The system-defined prefix matches the lower 2 bytes of the first MAC address in the burned-in MAC address pool that is programmed into the IDPROM. Use **connect fxos**, then **show module** to view the

MAC address pool. For example, if the range of MAC addresses shown for module 1 is b0aa.772f.f0b0 to b0aa.772f.f0bf, then the system prefix will be f0b0.

The user-defined prefix is an integer that is converted into hexadecimal. For an example of how the user-defined prefix is used, if you set a prefix of 77, then the chassis converts 77 into the hexadecimal value 004D (*yyxx*). When used in the MAC address, the prefix is reversed (*xxyy*) to match the chassis native form:

A2**4D.00***zz.zzzz*

For a prefix of 1009 (03F1), the MAC address is:

A2**F1.03***zz.zzzz*

Performance Scaling Factor for Multi-Instance Mode

The maximum throughput (connections, VPN sessions) for a platform is calculated for an appliance mode device's use of memory and CPU (and this value is shown in **show resource usage**). If you use multiple instances, then you need to calculate the throughput based on the percentage of CPU cores that you assign to the instance. For example, if you use an instance with 50% of the cores, then you should initially calculate 50% of the throughput. Moreover, the throughput available to an instance may be less than that available to an appliance.

For detailed instructions on calculating the throughput for instances, see https://www.cisco.com/c/en/us/products/collateral/security/firewalls/white-paper-c11-744750.html.

Instances and High Availability

You can use High Availability using an instance on 2 separate chassis; for example, if you have 2 chassis, each with 10 instances, you can create 10 High Availability pairs. You can also have standalone instances on the same chassis as High Availability instances. For detailed requirements, see Requirements and Prerequisites for Instances, on page 368.

Note Clustering is not supported.

Licenses for Instances

All licenses are consumed per chassis and not per instance. See the following details:

- Essentials licenses are assigned to the chassis as a whole, one per chassis.
- Feature licenses are assigned to each instance, but you only *consume* one license per feature per chassis.

Requirements and Prerequisites for Instances

Model Support

- Secure Firewall 3110

- Secure Firewall 3120
- Secure Firewall 3130
- Secure Firewall 3140
- Secure Firewall 4215
- Secure Firewall 4225
- Secure Firewall 4245

Note The Secure Firewall 3105 is not supported.

Maximum Container Instances and Resources per Model

For each container instance, you can specify the number of CPU cores (or more specifically, threads) to assign to the instance. We use the term "core" generically to account for different hardware architecture. RAM is dynamically allocated according to the number of cores, and disk space is set to 40 GB per instance.

Table 25: Maximum Container Instances and Resources per Model

Model	Max. Container Instances	Available CPU Cores (Threads)
Secure Firewall 3110	3	22
Secure Firewall 3120	5	30
Secure Firewall 3130	7	46
Secure Firewall 3140	10	62
Secure Firewall 4215	10	62
Secure Firewall 4225	15	126
Secure Firewall 4245	34	254

Software Requirements

- You can run different versions of Firewall Threat Defense software on each instance as long as they are all compatible with the version of FXOS running on the chassis.
- You cannot deploy an instance with a patch version of the Firewall Threat Defense software because it is not a complete bundle. You need to first deploy the major or maintenance version, and then patch it after deployment.

High Availability Requirements

- The two instances in a High Availability configuration must:
  - Be on separate chassis.

- Be on the same model.
- Have the same interfaces assigned. All interfaces must be preconfigured in the chassis identically before you enable High Availability.
- Use the same resource profile attributes. The profile names can be different, but the definitions need to match.

Firewall Management Center Requirements

For chassis management and all instances on the chassis, you must use the same Firewall Management Center due to the licensing implementation.

Guidelines and Limitations for Instances

General Guidelines

- A single Firewall Management Center must manage all instances on a chassis, as well as manage the chassis itself.
- For instances, the following features are not supported:
  - TLS crypto acceleration
  - Clustering
  - Firewall Management Center UCAPL/CC mode
  - Flow offload to hardware
- Primary management of the chassis by Security Cloud Control cloud-delivered Firewall Management Center and separate analytics-only management of the chassis by an on-prem Firewall Management Center is not supported. You can however add Security Cloud Control-managed instances to an analytics-only on-prem Firewall Management Center.

Management Interface

- No support for a data interface for chassis management; only the dedicated Management interface can be used
- No DHCP addressing for the Management interface

VLAN Subinterfaces

- This document discusses *chassis* VLAN subinterfaces only. You can separately create subinterfaces within the instance.
- If you assign a parent interface to an instance, it only passes untagged (non-VLAN) traffic. Do not assign the parent interface unless you intend to pass untagged traffic.
- Subinterfaces are supported on Data or Data-sharing type interfaces.

- You can create up to 500 VLAN IDs.
- You cannot use subinterfaces for an inline set or as a passive interface.
- If you use a subinterface for the failover link, then all subinterfaces on that parent, and the parent itself, are restricted for use as failover links. You cannot use some subinterfaces as failover links, and some as regular data interfaces.

EtherChannels

- You can configure up to 48 EtherChannels, limited by the number of physical interfaces.
- The EtherChannel can have up to 8 active interfaces.
- All interfaces in the EtherChannel must be the same media type and speed capacity. The media type can be either RJ-45 or SFP; SFPs of different types (copper and fiber) can be mixed. You cannot mix interface capacities (for example 1GB and 10GB interfaces) by setting the speed to be lower on the larger-capacity interface, unless you set the speed to **Detect SFP**; in this case, you can use different interface capacities, and the lowest common speed is used.
- The chassis does not support LACPDUs that are VLAN-tagged. If you enable native VLAN tagging on the neighboring switch using the Cisco IOS **vlan dot1Q tag native** command, then the chassis will drop the tagged LACPDUs. Be sure to disable native VLAN tagging on the neighboring switch.
- In Cisco IOS software versions earlier than 15.1(1)S2, the chassis did not support connecting an EtherChannel to a switch stack. With default switch settings, if the chassis EtherChannel is connected cross stack, and if the primary switch is powered down, then the EtherChannel connected to the remaining switch will not come up. To improve compatibility, set the **stack-mac persistent timer** command to a large enough value to account for reload time; for example, 8 minutes or 0 for indefinite. Or, you can upgrade to more a more stable switch software version, such as 15.1(1)S2.

Data-sharing Interfaces

- Maximum 14 instances per shared interface. For example, you can allocate Ethernet1/1 to Instance1 through Instance14.

  Maximum 10 shared interfaces per instance. For example, you can allocate Ethernet1/1.1 through Ethernet1/1.10 to Instance1.

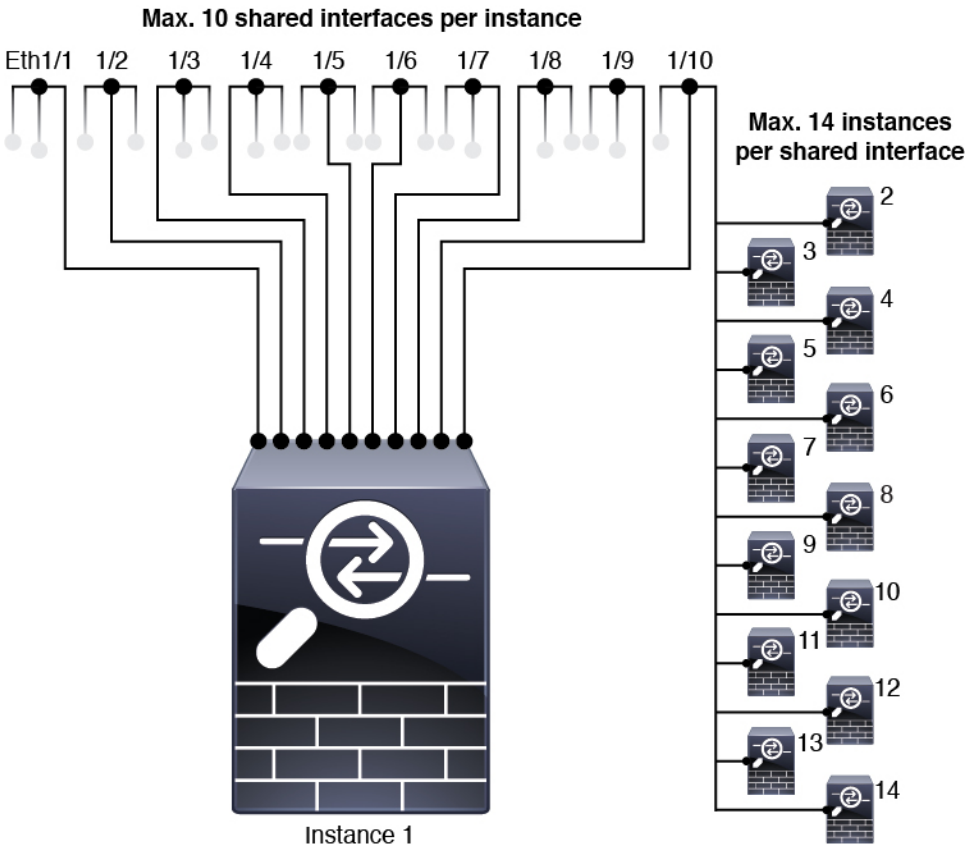

- You cannot use a data-sharing interface with a transparent firewall mode instance.
- You cannot use a data-sharing interface with inline sets or passive interfaces.
- You cannot use a data-sharing interface for the failover link.

Default MAC Addresses

- MAC addresses for all interfaces are taken from a MAC address pool. For subinterfaces, if you decide to manually configure MAC addresses, make sure you use unique MAC addresses for all subinterfaces on the same parent interface to ensure proper classification. See Automatic MAC Addresses for Instance Interfaces, on page 367.

Configure Instances

Before you configure instances, you need to enable multi-instance mode, add the chassis to the Firewall Management Center, and configure chassis interfaces. You can also customize the chassis settings.

Convert a Device to Multi-Instance Mode

Use this procedure to convert a device in the Firewall Management Center to multi-instance mode. The conversion takes approximately 15 minutes. The system reboots and, as part of changing the mode, erases

the configuration with the exception of the Management network settings and the admin password. The chassis hostname is set to "firepower-*model*." The Management IP address is assigned to the chassis for the management connection with the Firewall Management Center. When you add instances, they will use separate IP addresses on the Management interface and maintain their own management connections.

After you convert to multi-instance mode, you can use the Firewall Management Center to configure all chassis settings as well as instances. The Secure Firewall Chassis Manager or configuration at the FXOS CLI is not supported.

Note If you would prefer to use the CLI to convert to multi-instance mode before adding the device to the Firewall Management Center, see Enable Multi-Instance Mode at the CLI, on page 410.

Before you begin

Add the appliance-mode device to the Firewall Management Center as a standalone device. This device cannot use:

- A data interface for manager access
- DHCP for the Management interface
- Zero-Touch Provisioning

Procedure

Step 1 On **Devices** > **Device Management**, next to the device you want to convert to multi-instance mode, choose **More** (⋮) > **Convert to Multi-Instance**.

Figure 141: Convert to Multi-Instance

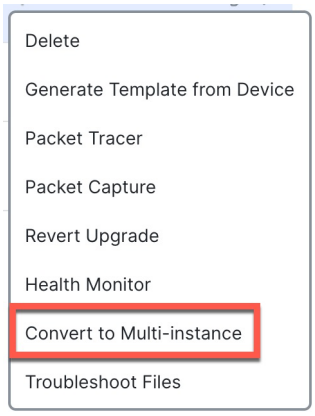

You can alternatively check the check-box next to multiple devices that you want to convert, and choose **Select Bulk Action** > **Convert to Multi-Instance**.

Figure 142: Bulk Conversion

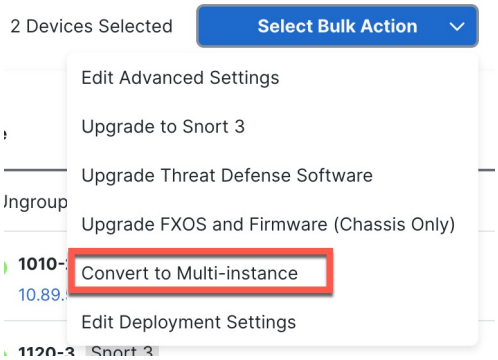

Step 2 Confirm that you want to perform the conversion and click **Continue**.

Figure 143: Conversion Confirmation

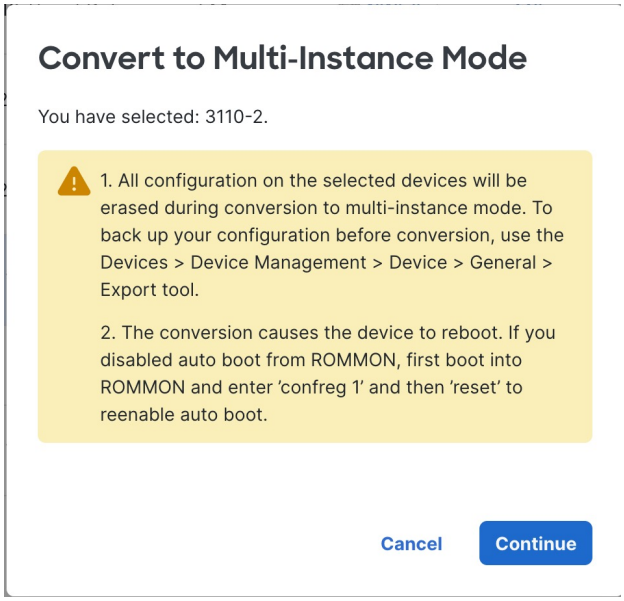

A readiness check is performed. The check might fail if, for example, a deployment is in progress.

Step 3 Optionally change the name of the chassis, and click **Convert to Multi-Instance**. By default, the device name is used, but you may want a different naming convention for multi-instance chassis.

Note
If the readiness check failed, for example because of a deployment in progress, you can wait for the process to complete and click to rerun the readiness check so you can continue.

Figure 144: Rename Chassis

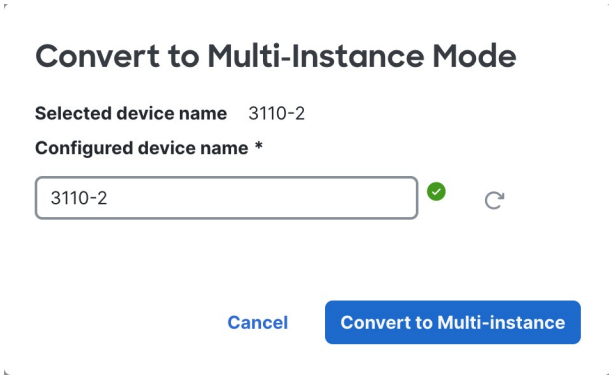

Wait for approximately 15 minutes, during which the device is removed from the device list, and then after conversion, it is re-added as a chassis.

Step 4 To view and configure the chassis, click **Manage** in the **Chassis** column, or click **Edit** (✎).

The **Chassis Manager** page opens for the chassis to the **Summary** page.

Figure 145: Chassis Summary

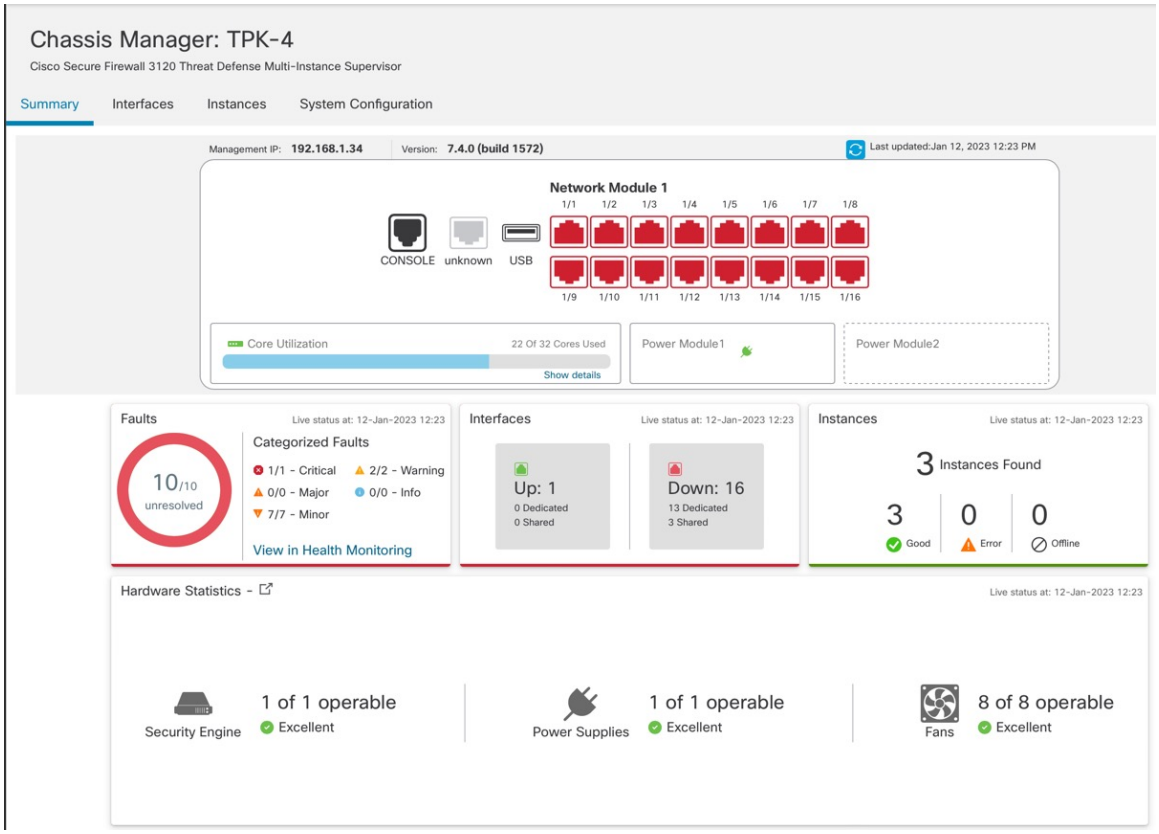

Configure Chassis Interfaces

At the chassis-level, you configure basic Ethernet settings of physical interfaces, VLAN subinterfaces for instances, and EtherChannel interfaces. By default, physical interfaces are disabled.

Note To configure breakout ports and perform other network module operations, see Manage the Network Module for the Secure Firewall 3100/4200, on page 790.

Note For information about the **Sync Device** button, see Sync Interface Changes with the Firewall Management Center, on page 789.

Configure a Physical Interface

You can physically enable and disable interfaces, as well as set the interface speed and duplex and other hardware settings. To use an interface, it must be physically enabled for the chassis and logically enabled in the instance. By default, physical interfaces are disabled. For VLAN subinterfaces, the admin state is inherited from the parent interface.

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column or click **Edit** ( ).

Figure 146: Manage Chassis

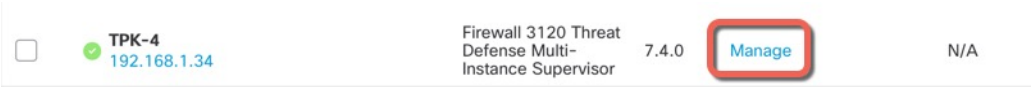

The **Chassis Manager** page opens for the chassis to the **Summary** page.

Step 2 Click **Interfaces**.

Figure 147: Interfaces

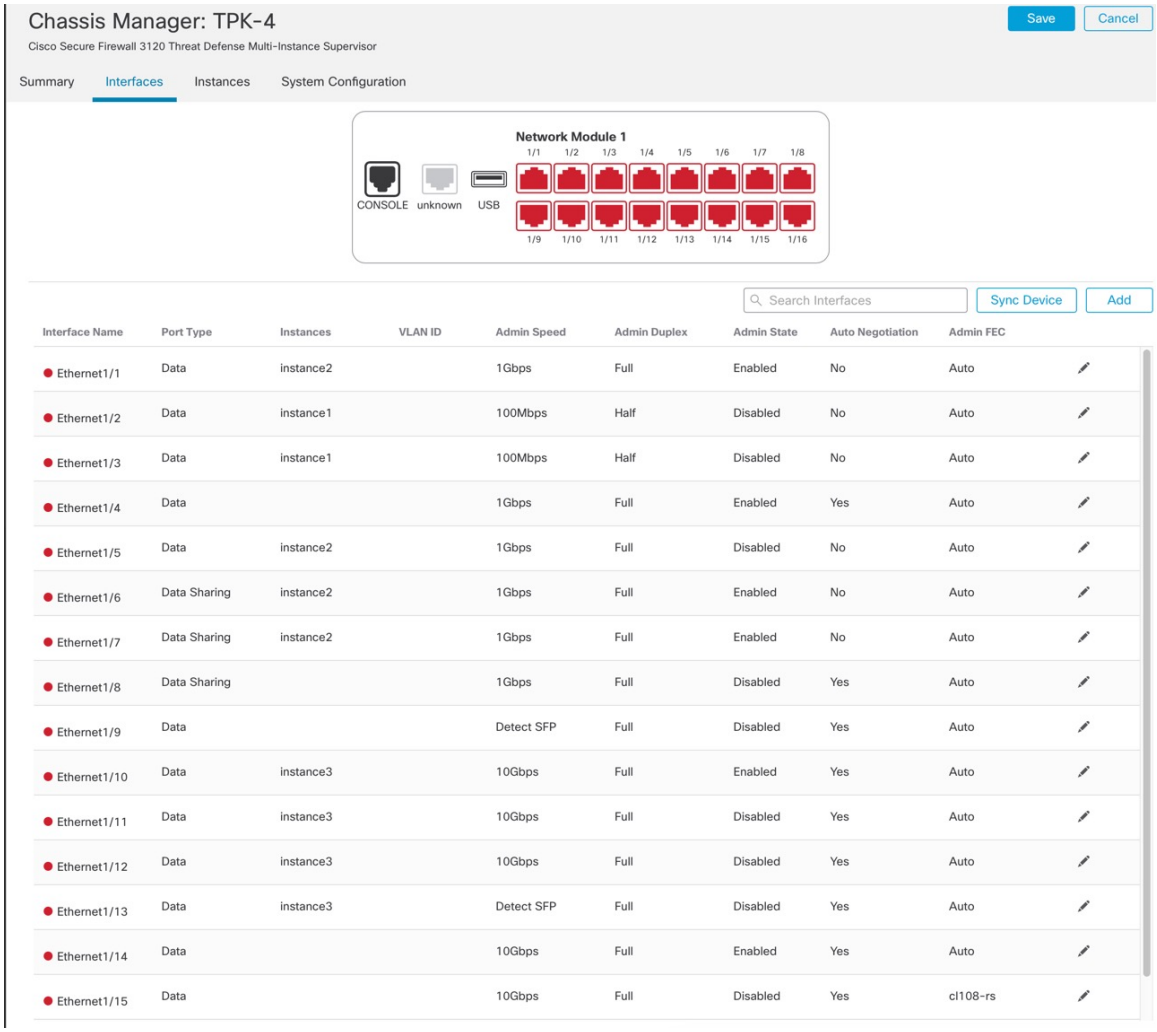

Step 3 Click **Edit** (✎) for the interface you want to edit.

Figure 148: Edit Physical Interface

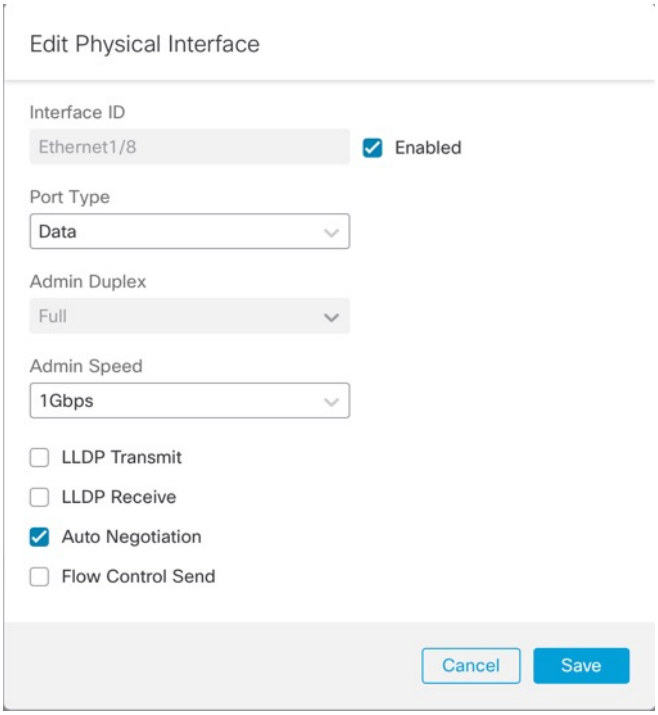

Step 4 Enable the interface by checking the **Enabled** check box.

Step 5 For the **Port Type**, choose **Data** or **Data Sharing**.

Figure 149: Port Type

Step 6 Set the **Admin Duplex**.

Speeds of 1Gbps and higher only support **Full** duplex. SFP interfaces only support **Full** duplex.

Step 7 Set the **Admin Speed**.

For SFPs, Choose **Detect SFP** to detect the speed of the installed SFP module and use the appropriate speed. Duplex is always Full, and auto-negotiation is always enabled. This option is useful if you later change the network module to a different model, and want the speed to update automatically.

Step 8 (Optional) Check **LLDP Transmit** and/or **LLDP Receive** to enable Link Layer Discovery Protocol (LLDP) packets.

Step 9 (Optional) Check **Flow Control Send** to enable pause (XOFF) frames for flow control.

Flow control enables connected Ethernet ports to control traffic rates during congestion by allowing congested nodes to pause link operation at the other end. If the threat defense port experiences congestion (exhaustion of queuing resources on the internal switch) and cannot receive any more traffic, it notifies the other port by

sending a pause frame to stop sending until the condition clears. Upon receipt of a pause frame, the sending device stops sending any data packets, which prevents any loss of data packets during the congestion period.

Note
The Firewall Threat Defense supports transmitting pause frames so that the remote peer can rate-control the traffic.

However, receiving of pause frames is not supported.

The internal switch has a global pool of 8000 buffers of 250 bytes each, and the switch allocates buffers dynamically to each port. A pause frame is sent out every interface with flow control enabled when the buffer usage exceeds the global high-water mark (2 MB (8000 buffers)); and a pause frame is sent out of a particular interface when its buffer exceeds the port high-water mark (.3125 MB (1250 buffers)). After a pause is sent, an XON frame can be sent when the buffer usage is reduced below the low-water mark (1.25 MB globally (5000 buffers); .25 MB per port (1000 buffers)). The link partner can resume traffic after receiving an XON frame.

Only flow control frames defined in 802.3x are supported. Priority-based flow control is not supported.

Step 10 (Optional) Check **Auto Negotiation** to set the interface to negotiate the speed, link status, and flow control. You cannot edit this setting for speeds lower than 1Gbps. For SFP interfaces, you can only disable auto-negotiation when the speed is set to 1Gbps.

Step 11 Click **Save** and then **Save** in the top right of the **Interfaces** page.

You can now **Deploy** the policy to the chassis. The changes are not active until you deploy them.

Configure an EtherChannel

An EtherChannel (also known as a port channel) can include up to 8 member interfaces of the same media type and capacity, and must be set to the same speed and duplex. The media type can be either RJ-45 or SFP; SFPs of different types (copper and fiber) can be mixed. You cannot mix interface capacities (for example 1GB and 10GB interfaces) by setting the speed to be lower on the larger-capacity interface, unless you set the speed to **Detect SFP**; in this case, you can use different interface capacities, and the lowest common speed is used.

The Link Aggregation Control Protocol (LACP) aggregates interfaces by exchanging the Link Aggregation Control Protocol Data Units (LACPDUs) between two network devices. LACP coordinates the automatic addition and deletion of links to the EtherChannel without user intervention. It also handles misconfigurations and checks that both ends of member interfaces are connected to the correct channel group. “On” mode cannot use standby interfaces in the channel group when an interface goes down, and the connectivity and configurations are not checked.

When the chassis creates an EtherChannel, the EtherChannel stays in a **Suspended** state for Active LACP mode or a **Down** state for On LACP mode until you assign it to a logical device, even if the physical link is up. The EtherChannel will be brought out of this **Suspended** state when it is added to an instance.

Before you begin

Enable physical interfaces and set hardware parameters. See Configure a Physical Interface, on page 376.

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column or click **Edit** (✎).

Figure 150: Manage Chassis

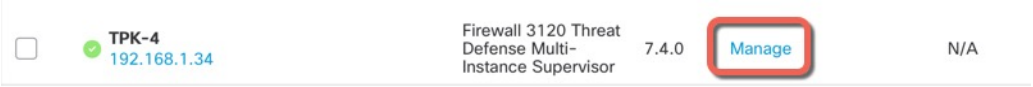

The **Chassis Manager** page opens for the chassis to the **Summary** page.

Step 2 Click **Interfaces**.

Figure 151: Interfaces

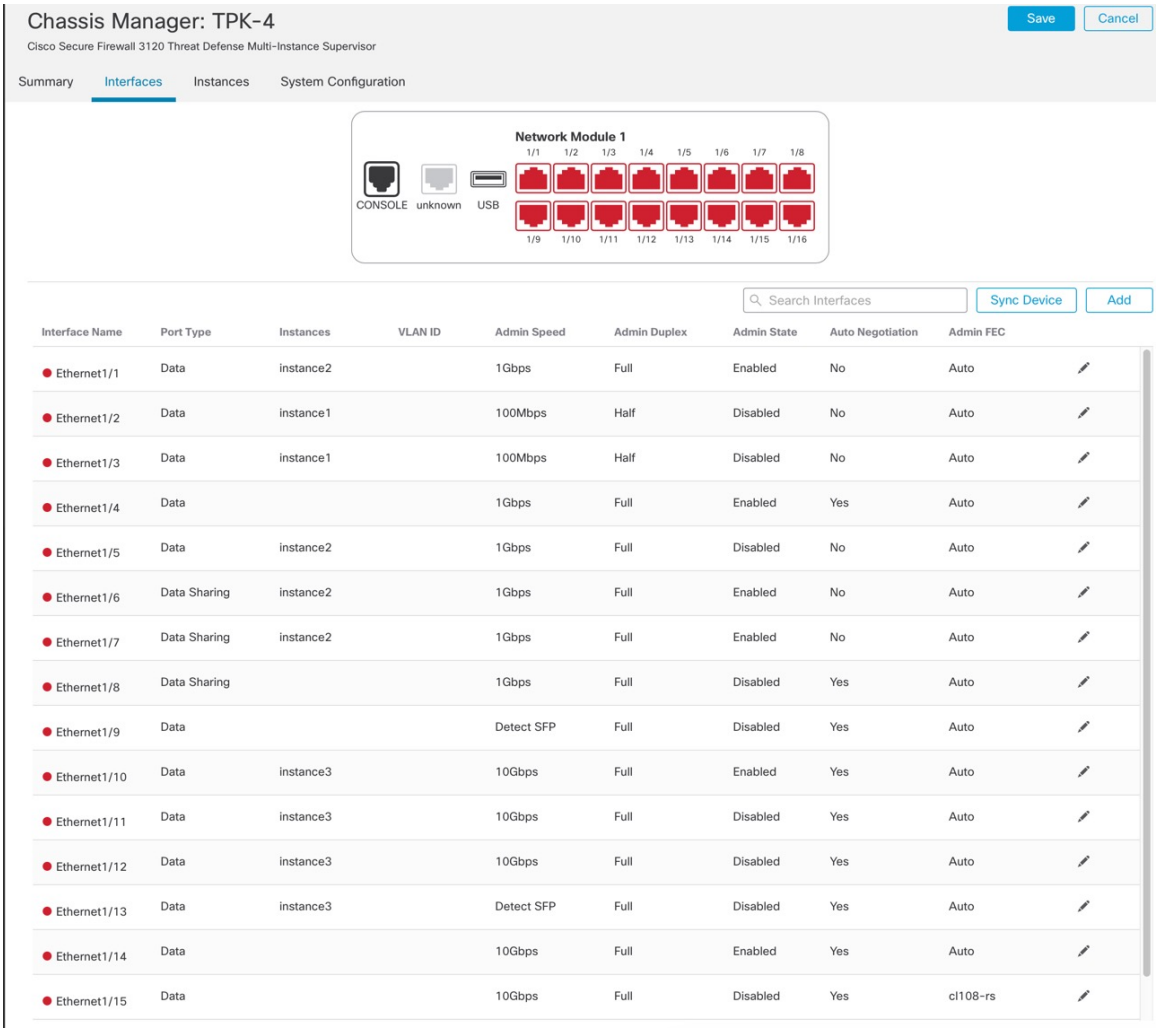

Step 3 Click **Add** > **EtherChannel Interface**.

Figure 152: Add EtherChannel

Step 4 Set the following **Interfaces** parameters.

Figure 153: Interfaces Settings

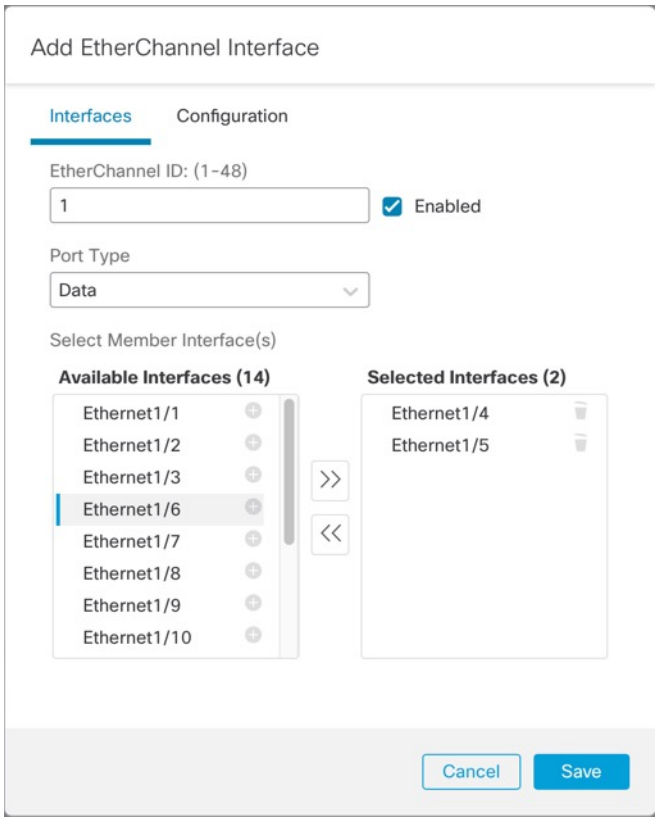

a) For the **EtherChannel ID**, specify an ID between 1 and 48.
b) Check **Enabled**.
c) For the **Port Type**, choose **Data** or **Data Shared**.

   For information about the port type, see Interface Types, on page 357.

d) To add a physical interface to the EtherChannel, select click **Add** ( ) in the **Available Interfaces** list to move it to the **Selected Interfaces** list.

   To add or remove all interfaces, click the double arrow button.

   Note
   You cannot add an interface that is already assigned to an instance.

Step 5 (Optional) Set the following **Configuration** parameters.

Many of these settings (excluding the LACP settings) set the requirements for interfaces to be included in the EtherChannel; they do not override the settings of member interfaces. So if you check **LLDP Transmit**, for example, you should only add interfaces that have that setting. If you set the **Admin Speed** to 1Gbps, then only 1Gbps interfaces can be included.

Figure 154: Configuration Settings

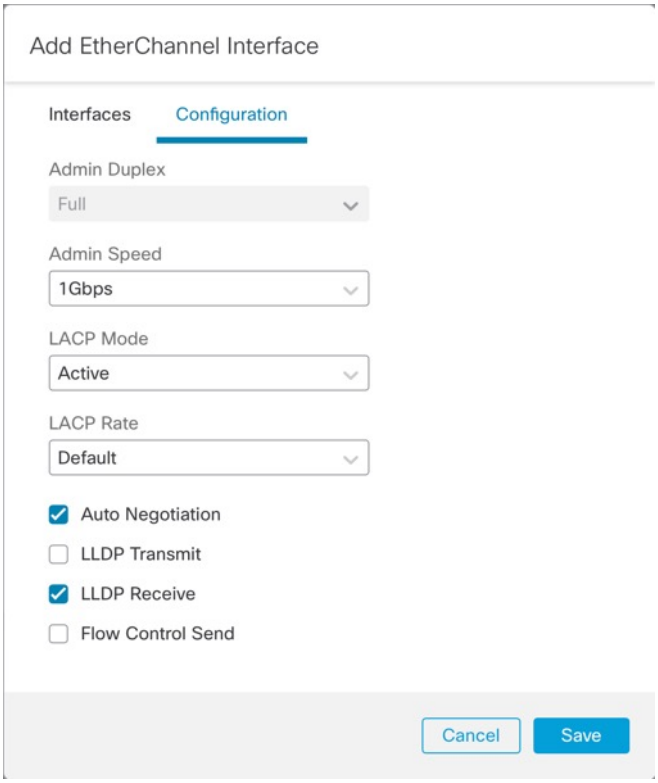

a) Choose the required **Admin Duplex** for the member interfaces, **Full Duplex** or **Half Duplex**.

If you add a member interface that is configured with the specified duplex, it will not successfully join the port channel.

b) Choose the required **Admin Speed** for the member interfaces from the drop-down list.

If you add a member interface that is not at the specified speed, it will not successfully join the port channel.

c) Choose the **LACP Mode**, **Active** or **On**.

- **Active**—Sends and receives LACP updates. An active EtherChannel can establish connectivity with either an active or a passive EtherChannel. You should use the active mode unless you need to minimize the amount of LACP traffic.
- **On**—The EtherChannel is always on, and LACP is not used. An "on" EtherChannel can only establish a connection with another "on" EtherChannel.

Note
It may take up to three minutes for an EtherChannel to come up to an operational state if you change its mode from On to Active or from Active to On.

d) Choose the **LACP Rate**, **Default**, **Fast**, or **Normal**.

The default is **Fast**.

e) Choose the required Link Layer Discovery Protocol (LLDP) settings for member interfaces by checking **LLDP Transmit** and/or **LLDP Receive**.

f) Check the required **Flow Control Send** setting for member interfaces.

Step 6 Click **Save** and then **Save** in the top right of the **Interfaces** page.

You can now **Deploy** the policy to the chassis. The changes are not active until you deploy them.

Configure a Subinterface

You can add up to 500 subinterfaces to your chassis.

VLAN IDs per interface must be unique, and within an instance, VLAN IDs must be unique across all assigned interfaces. You can reuse VLAN IDs on *separate* interfaces as long as they are assigned to different instances. However, each subinterface still counts towards the limit even though it uses the same ID.

This section discusses *FXOS* VLAN subinterfaces only. You can separately create subinterfaces within the instance. See Chassis Interfaces vs. Instance Interfaces, on page 357.

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column or click **Edit** ( ).

Figure 155: Manage Chassis

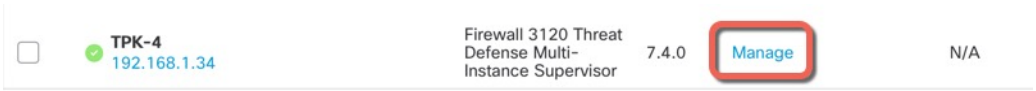

The **Chassis Manager** page opens for the chassis to the **Summary** page.

Step 2 Click **Interfaces**.

Figure 156: Interfaces

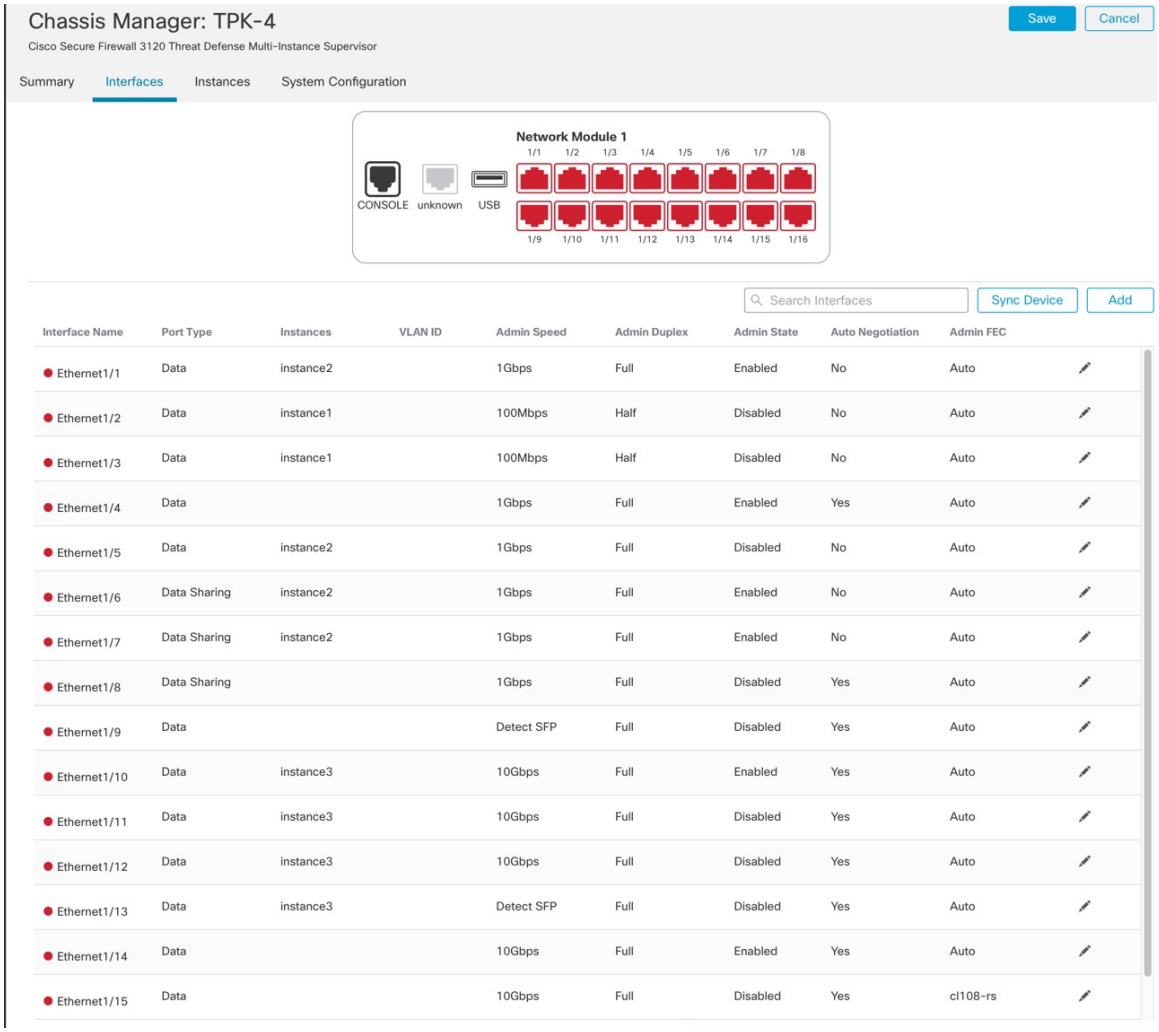

Step 3 Click **Add** > **Subinterface**.

Figure 157: Add Subinterface

Step 4 Set the following parameters.

Figure 158: Subinterface Settings

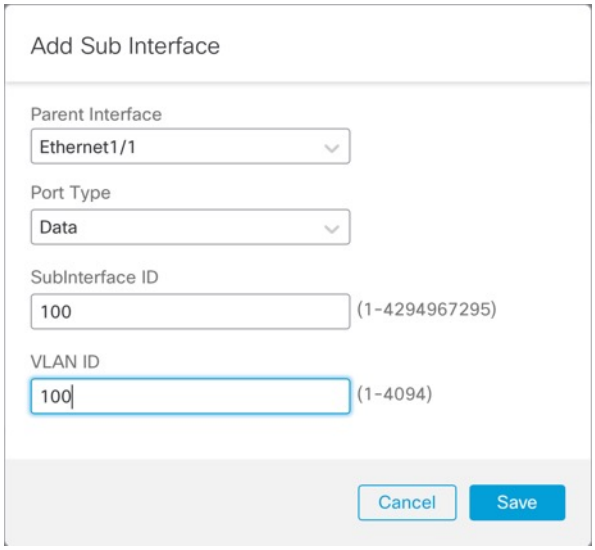

a)

Step 5 Click **Save** and then **Save** in the top right of the **Interfaces** page.

You can now **Deploy** the policy to the chassis. The changes are not active until you deploy them.

Add an Instance

You can add one or more instances to a chassis in multi-instance mode. The number of supported instances depends on your model; see Requirements and Prerequisites for Instances, on page 368.

Before you begin

Convert a Device to Multi-Instance Mode, on page 372.

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column or click **Edit** ( ).

Figure 159: Manage Chassis

The **Chassis Manager** page opens for the chassis to the **Summary** page.

Step 2 Click **Instances**, and click **Add Instance**.

Figure 160: Instances

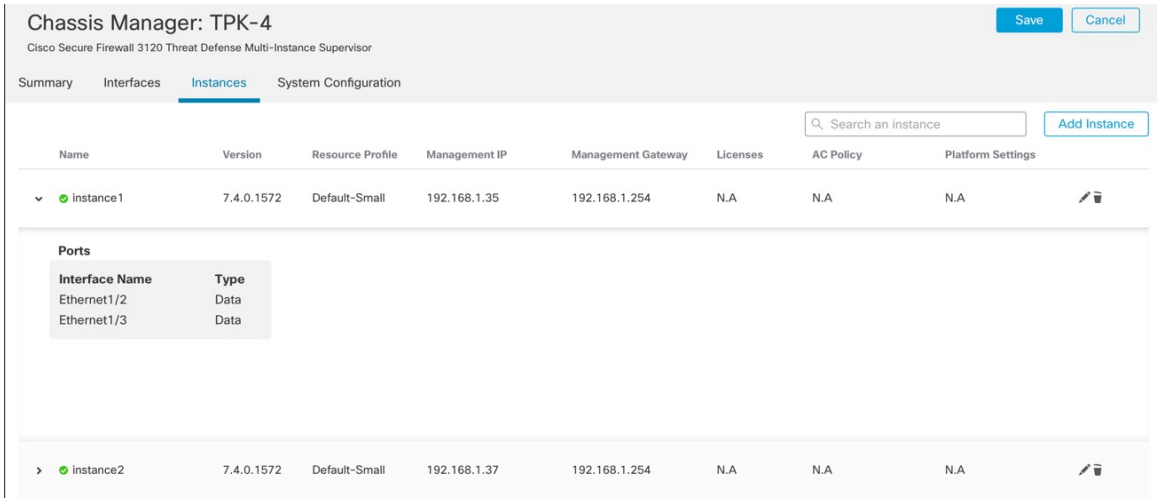

Step 3 On **Agreement**, check **I understand and accept the agreement**, then click **Next**.

Figure 161: Agreement

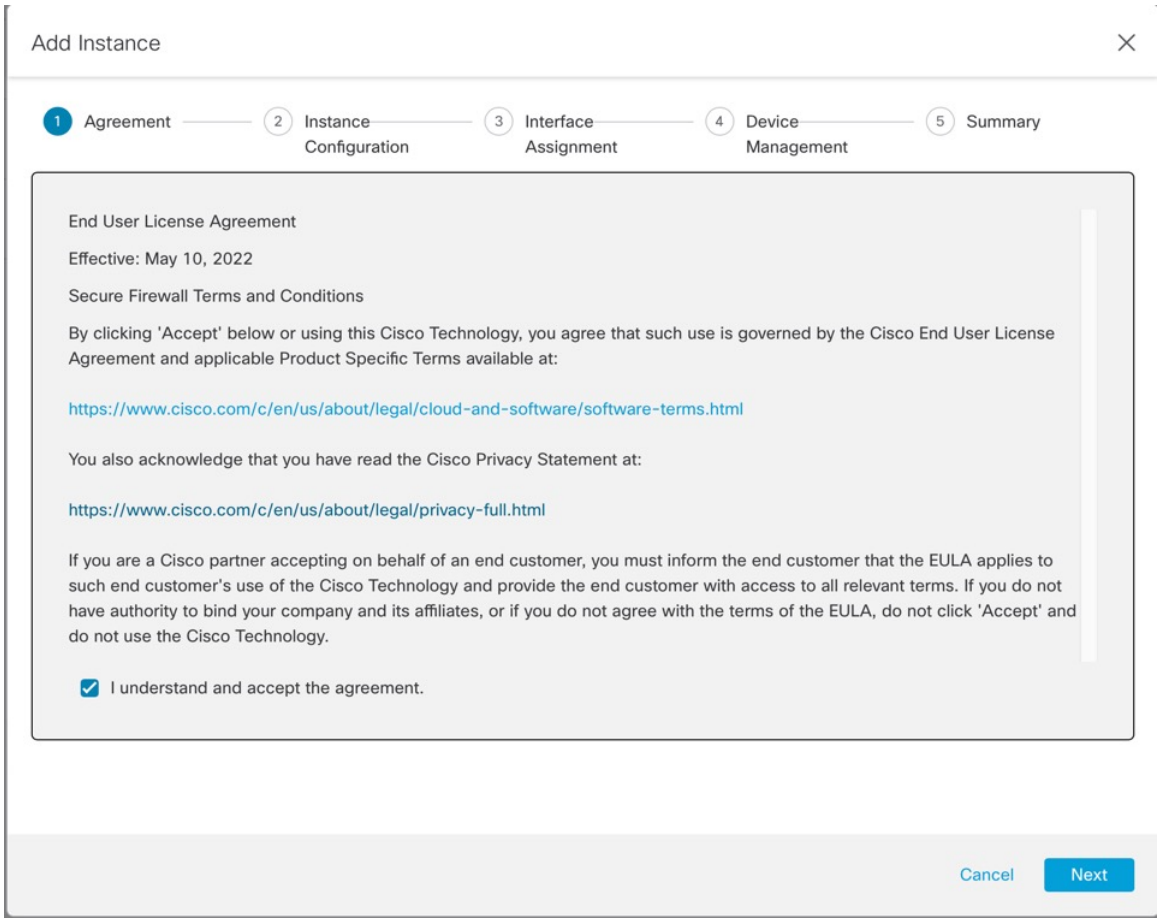

Step 4 On **Instance Configuration**, set the instance parameters, then click **Next**.

Figure 162: Instance Configuration

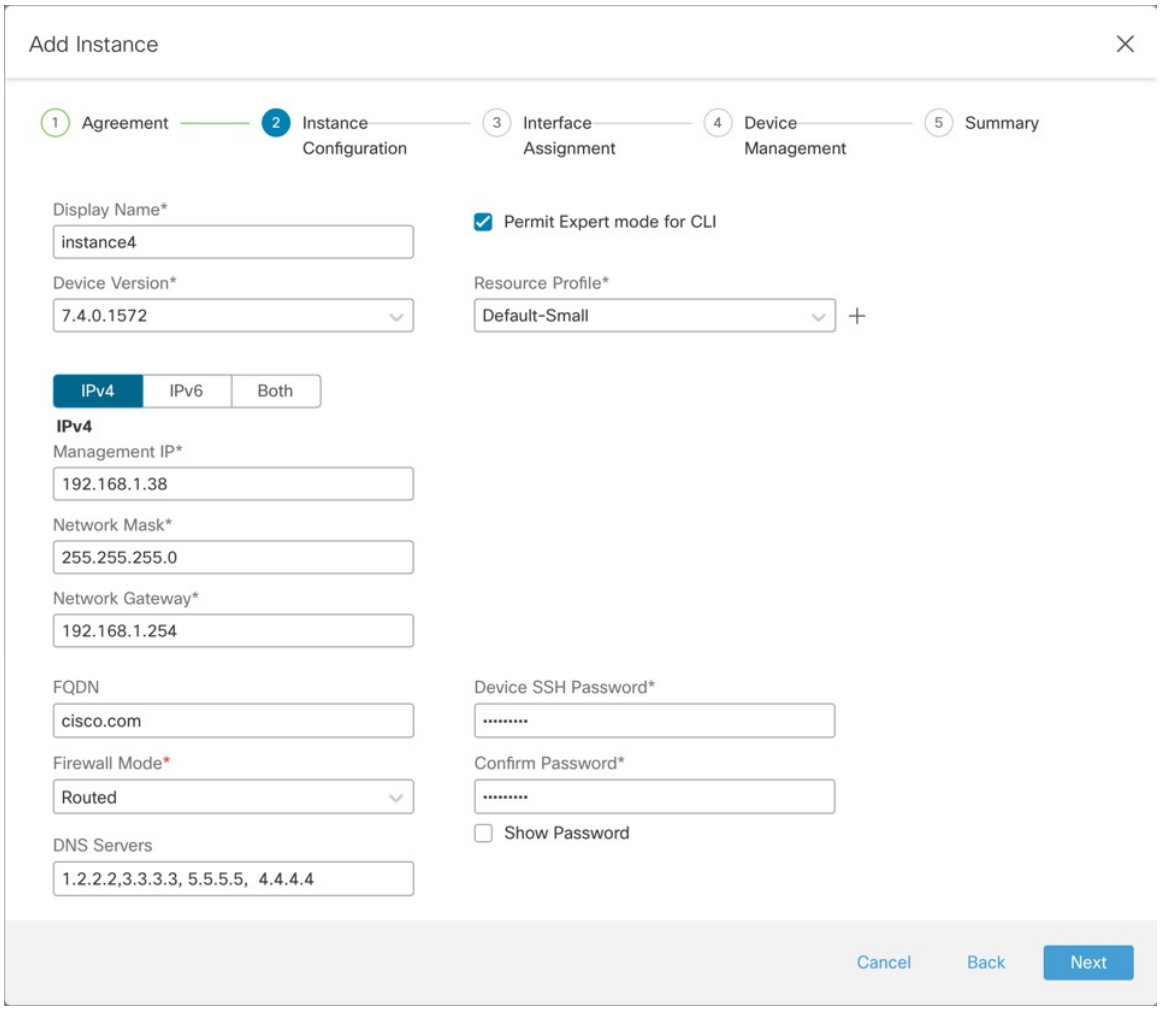

- **Display Name**
- **Device Version**—Versions listed are packages currently downloaded to the chassis. Patch versions are not listed and cannot be used because they don't contain the entire bundle. To upgrade to a new package, see **Devices** > **Upgrade** > **Chassis Upgrade**. When you upgrade, both the old Firewall Threat Defense version and the new Firewall Threat Defense version will be listed in the menu. To download an older package, you need to use the FXOS CLI. Note: Both FXOS and Firewall Threat Defense images are included in the same package. See the troubleshooting guide for more information.

  For example:

```
firepower-3110# scope firmware
firepower-3110# download image
https://10.10.7.89/Cisco_FTD_SSP_FP3K_Upgrade-7.4.1-1.sh.DEV.tar
Please use the command 'show download-task' or 'show download-task detail' to check
download progress.
% Download-task Cisco_FTD_SSP_FP3K_Upgrade-7.4.1-1.sh.DEV.tar : completed successfully.
```

- **IPv4**, **IPv6**, or **Both**—Set a **Management IP** address on the same network as the chassis Management interface. Set the **Network Mask** and gateway (likely the same gateway as the chassis). The chassis Management interface is shared with each instance, and each instance has its own IP address on the network. You can SSH to this IP address by default to reach the Firewall Threat Defense CLI.
- (Optional) **FQDN**
- **Firewall Mode**—**Routed** or **Transparent**. For more information about the firewall mode, see Transparent or Routed Firewall Mode, on page 285.
- **DNS Servers**—Enter a comma-separated list of DNS servers for management traffic only.
- (Optional) **Permit Expert Mode for CLI**—Expert Mode provides Firewall Threat Defense shell access for advanced troubleshooting.

  If you enable this option, then users who access the instance directly from an SSH session can enter Expert Mode. If you disable this option, then only users who access the instance from the FXOS CLI can enter Expert Mode. We recommend disabling this option to increase isolation between instances.

  Use Expert Mode only if a documented procedure tells you it is required, or if the Cisco Technical Assistance Center asks you to use it. To enter this mode, use the **expert** command in the Firewall Threat Defense CLI.
- **Resource Profile**—The resource profile sets the number of CPU cores; RAM is dynamically allocated according to the number of cores, and disk space is set to 40 GB per instance. The chassis includes the following default resource profiles: Default-Small, Default-Medium, and Default-Large. You can add additional profiles for this chassis by clicking **Add** (+). You cannot later edit the resource profile.

Figure 163: Add Resource Profile

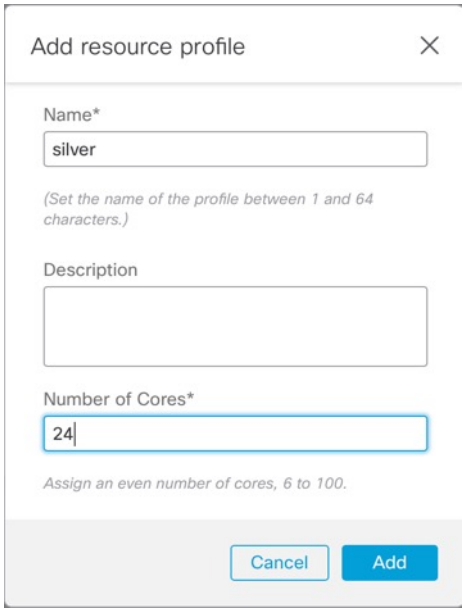

  - The minimum number of cores is 6.

  Note

Instances with a smaller number of cores might experience relatively higher CPU utilization than those with larger numbers of cores. Instances with a smaller number of cores are more sensitive to traffic load changes. If you experience traffic drops, try assigning more cores.

- You can assign cores as an even number (6, 8, 10, 12, 14 etc.) up to the maximum.
- The maximum number of cores available depends on the model; see Requirements and Prerequisites for Instances, on page 368.

If you later assign a different resource profile, then the instance will reload, which can take approximately 5 minutes. Note that for an established High Availability pair, if you assign a different-sized resource profile, be sure to make all members the same size as soon as possible.

- **Device SSH Password**—Set the Firewall Threat Defense admin user password for CLI access, either SSH or console. Repeat the password in the **Confirm Password** field.

Step 5 On **Interface Assignment**, assign the chassis interfaces to the instance, then click **Next**.

Figure 164: Interface Assignment

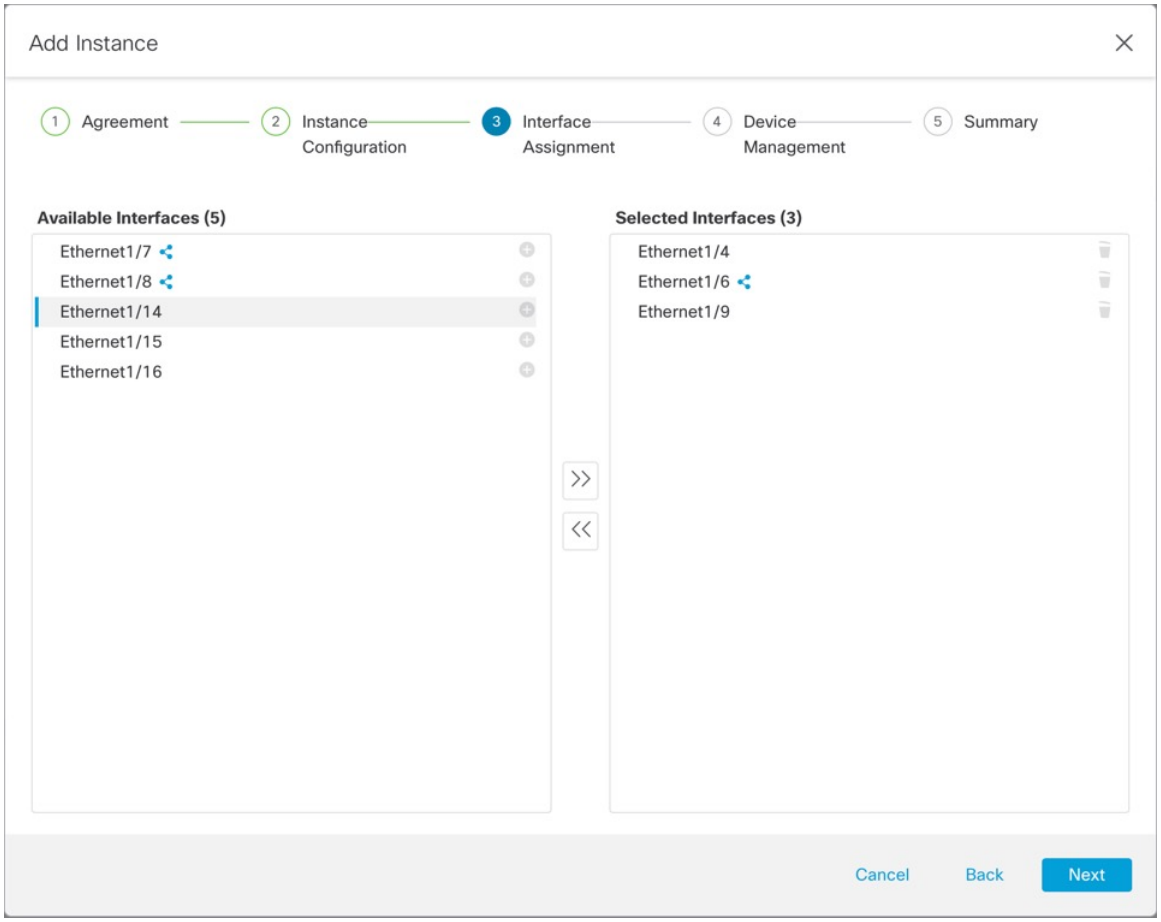

Shared interfaces show the sharing icon ( ).

Step 6 On **Device Management**, set the device-specific settings, then click **Next**.

Figure 165: Device Management

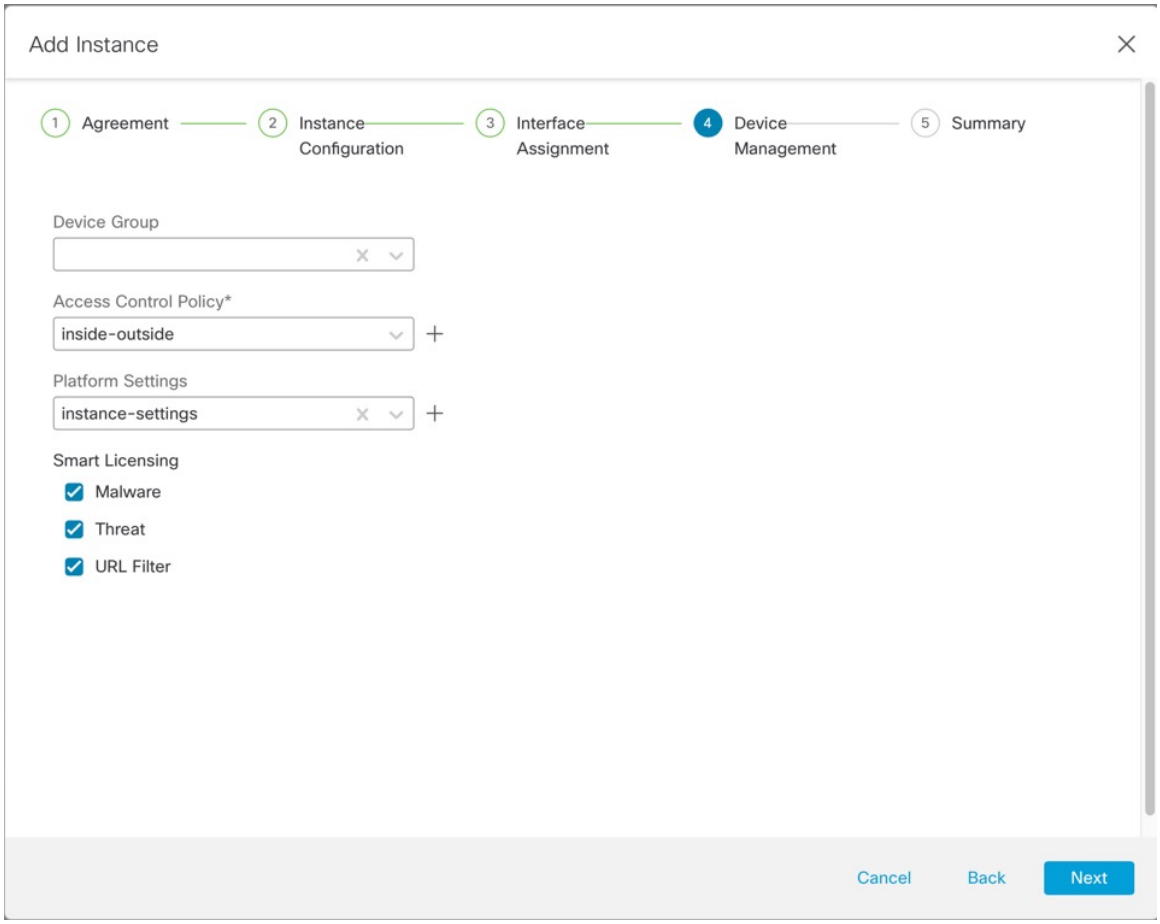

- **Device Group**
- **Access Control Policy**—Choose an existing access control policy, or create a new policy.
- **Platform Settings**—Choose an existing platform setting policy, or create a new policy.
- **Smart Licensing**

Step 7 On **Summary**, confirm your settings, then click **Save**.

Figure 166: Summary

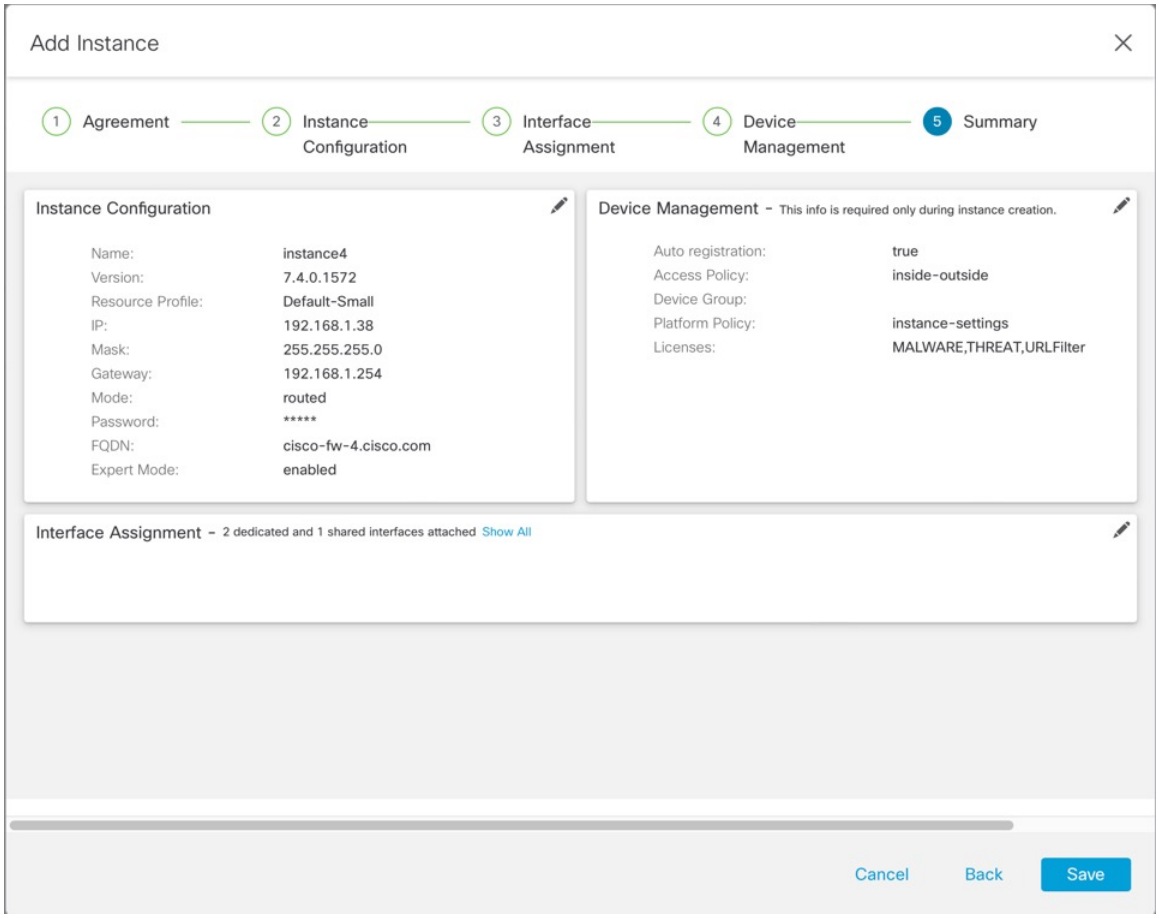

You can edit any settings on this screen before saving the instance. After you save, the instance is added to the **Instances** screen.

Step 8 On the **Instances** screen, click **Save**.

Step 9 Deploy the chassis configuration.

After deployment, the instance will be added as a device on the **Device Management** page.

Customize the System Configuration

You can configure chassis-level settings such as SNMP. You can also import or export the chassis FXOS configuration.

Configure SNMP

You can access chassis-level MIBs through the data interface of one of the instances, which you specify in the chassis system configuration. You can only use this instance for chassis SNMP information. You cannot access SNMP through the chassis Management interface.

Before you begin

Configure SNMP for one of the instances. See SNMP, on page 966.

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column or click **Edit** ( ).

Figure 167: Manage Chassis

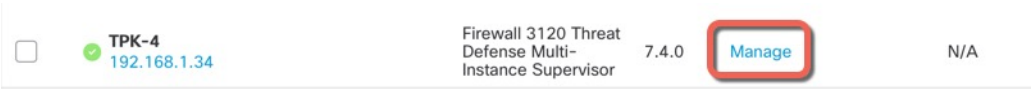

The **Chassis Manager** page opens for the chassis to the **Summary** page.

Step 2 Click **System Configuration**.

Step 3 Click **SNMP**, and choose the instance from the drop-down list.

Figure 168: SNMP

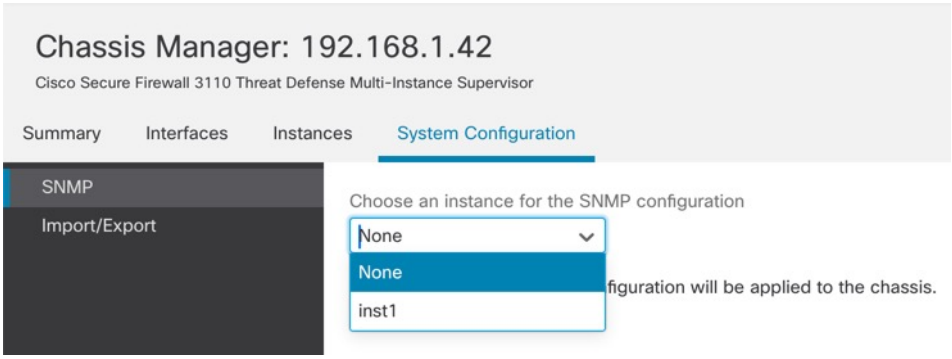

SNMP on the chassis will be accessible from the selected instance.

Step 4 Click **Save**.

Step 5 Deploy the chassis configuration.

Import or Export the Chassis Configuration

You can use the configuration export feature to export an XML file containing chassis configuration settings to your local computer. You can later import that configuration file to quickly apply the configuration settings to your chassis to return to a known good configuration or to recover from a system failure. You can also import the chassis configuration to a new chassis, for example for an RMA, as long as the prerequisites are met.

When exporting, only chassis configurations are exported; instance configuration settings are not exported. Instances need to be backed up separately using the device backup/restore feature.

When importing, all existing configurations on the chassis will be replaced by the configuration in the import file.

Before you begin

For the chassis where you want to import a configuration, the following characteristics must match:

- Same chassis software version
- Same Firewall Threat Defense instance images
- Same network modules

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column or click **Edit** ( ).

Figure 169: Manage Chassis

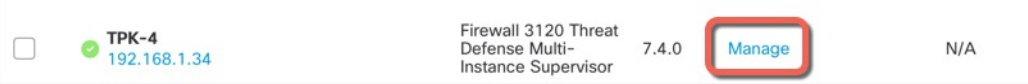

The **Chassis Manager** page opens for the chassis to the **Summary** page.

Step 2 Click **System Configuration**.

Step 3 Click **Import/Export**.

Step 4 To export the configuration, follow these steps.

a) In the **Export** area, click **Click here to export**.

Figure 170: Create Export File

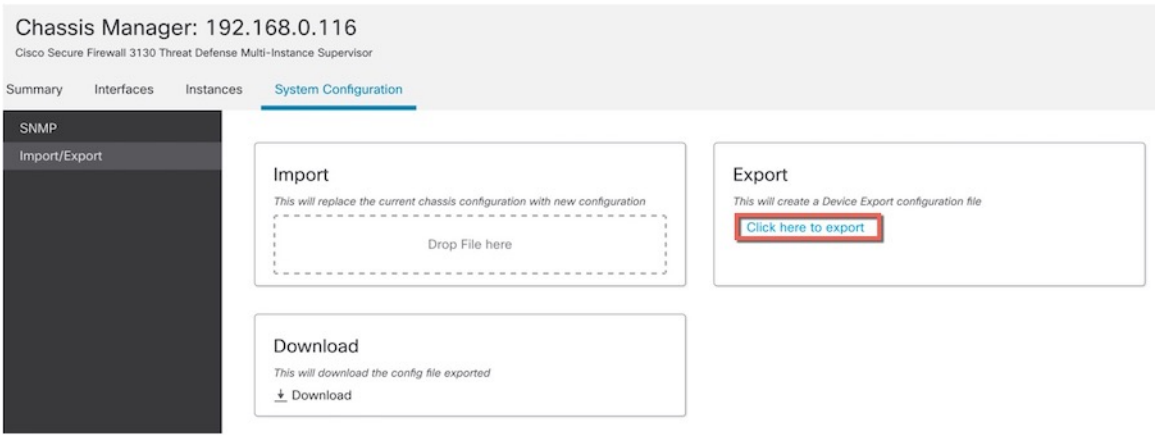

b) Monitor the notifications for the **Export file created successfully** message.

Figure 171: Export File Created Successfully

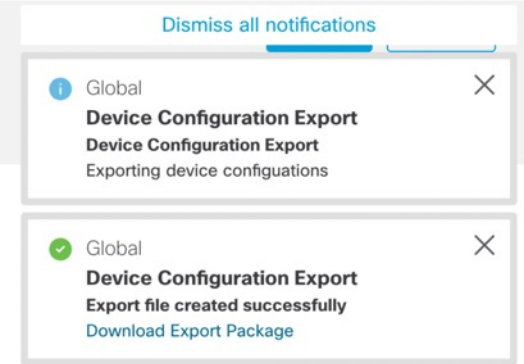

c) Download the export file by clicking the notification message (**Download Export Package**) or by clicking **Download**.

Figure 172: Download

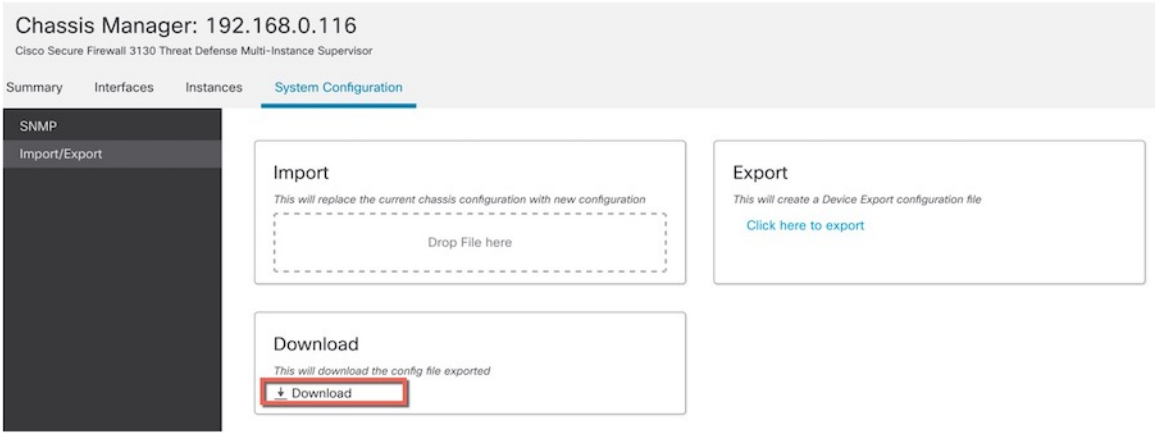

The file is saved with the **.sfo** extension.

Step 5 To import a configuration, drag the **.sfo** file on the **Import** > **Drop File here** area.

Figure 173: Import

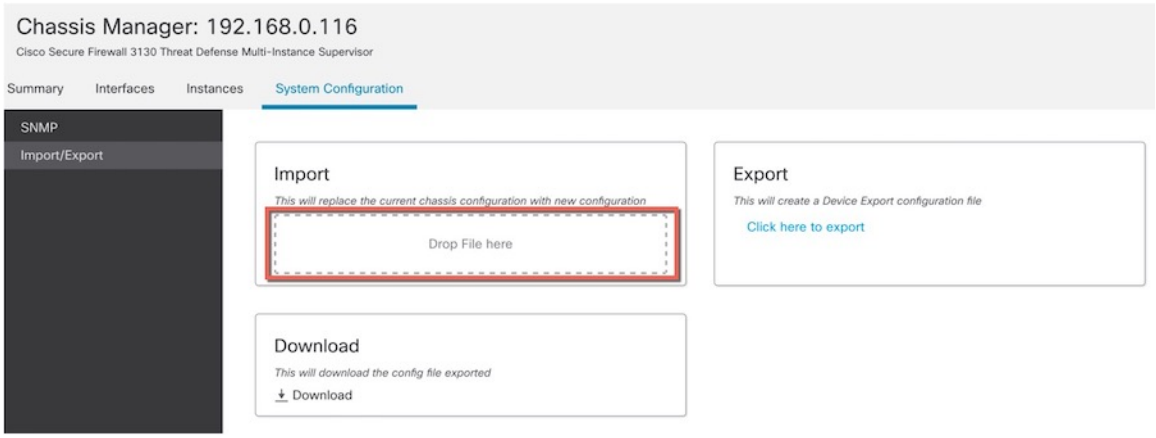

Configure Chassis Platform Settings

Chassis platform settings configure a range of features for managing the chassis. You can share the policy among multiple chassis. If you want different settings per chassis, you must create multiple policies.

Create a Chassis Platform Settings Policy

Use the **Platform Settings** page (**Devices** > **Platform Settings**) to manage platform settings policies. This page indicates the type of device for each policy. The **Status** column shows the device targets for the policy.

Procedure

Step 1 Choose **Devices** > **Platform Settings**.

Step 2 For an existing policy, you can **Copy** (), **Edit** (), or **Delete** () the policy.

Caution
You should not delete a policy that is the last-deployed policy on any of its target devices, even if it is out of date. Before you delete the policy completely, it is good practice to deploy a different policy to those targets.

Step 3 To create a new policy, click **New Policy**.

a) Choose **Chassis Platform Settings** from the drop-down list.
b) Enter a **Name** for the new policy and optionally, a **Description**.
c) Optionally, choose the **Available Chassis** where you want to apply the policy and click **Add** (or drag and drop) to add the selected chassis. You can enter a search string in the **Search** field to narrow the list of chassis.
d) Click **Save**.

The system creates the policy and opens it for editing.

Step 4 To change the target chassis for a policy, click **Edit** (✎) next to the platform settings policy that you want to edit.

a) Click **Policy Assignment**.
b) To assign a chassis to the policy, select it in the **Available Chassis** list and click **Add**. You can also drag and drop.
c) To remove a chassis assignment, click **Delete** (🗑) next to a chassis in the **Selected Chassis** list.
d) Click **OK**.

Configure DNS

You need to specify a DNS server if the chassis requires resolution of hostnames to IP addresses. These chassis DNS settings are separate from the DNS settings per instance, which are configured in the device platform settings.

When you configure multiple DNS servers, the chassis uses servers in a random order. You can configure up to four servers across four DNS server groups. For example, you can configure a single server group with four servers, or you can configure four server groups with one server each.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the chassis policy.

Step 2 Choose **DNS**.

Figure 174: DNS

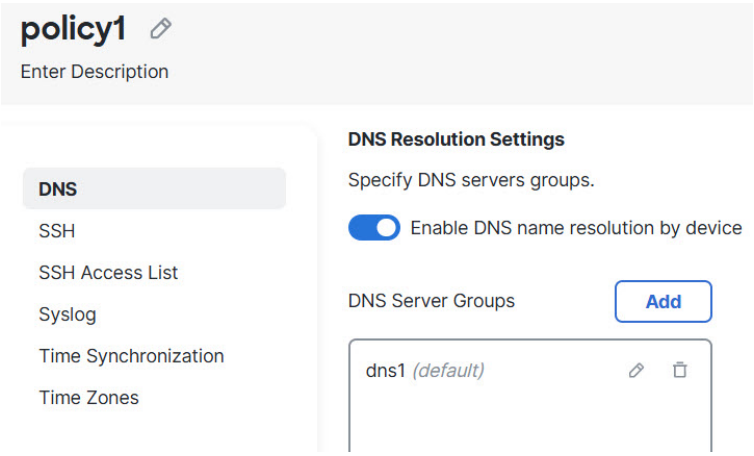

Step 3 Enable the **Enable DNS name resolution by device** slider.

Step 4 Click **Add** to add a DNS server group.

Figure 175: Add DNS Server Group

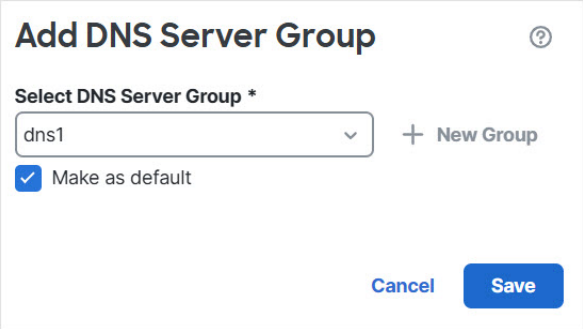

Step 5 Either select an existing DNS server group (see Creating DNS Server Group Objects, on page 1384), or click (+) **New Group**.

If you add a new group, you see the following dialog box. Provide a name and up to four DNS server IP addresses as comma-separated values, and click **Add**.

Figure 176: New DNS Server Group Object

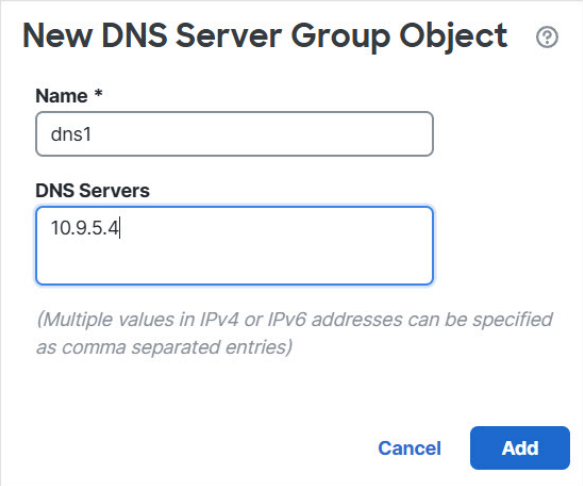

Step 6 Click **Save** to add the DNS server to the list.

Step 7 Repeat these steps to add additional server groups.

Make sure you only identify a maximum of four DNS servers in all groups combined.

Step 8 Click **Save** to save all policy changes.

Configure SSH and SSH Access List

To allow SSH sessions from the admin user to the chassis on the Management interface, enable the SSH server and configure the allowed networks.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the chassis policy.

Step 2 Choose **SSH**.

Step 3 To enable SSH access to the chassis, enable the **Enable SSH Server** slider.

Figure 177: SSH

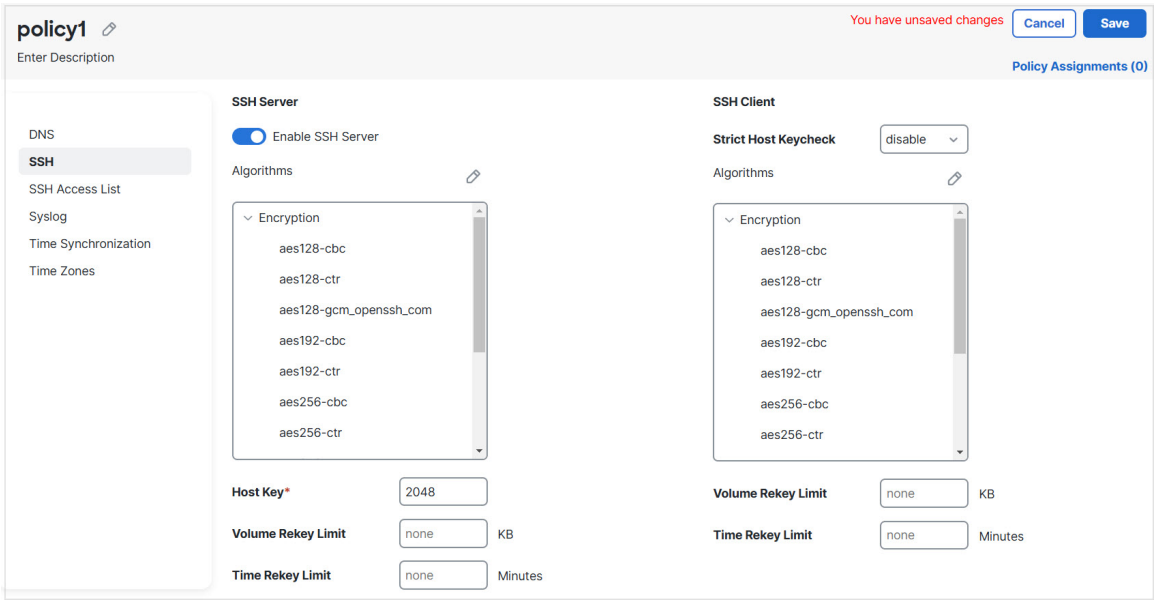

Step 4 To set the allowed **Algorithms**, click **Edit** ( ).

Figure 178: Add Algorithms

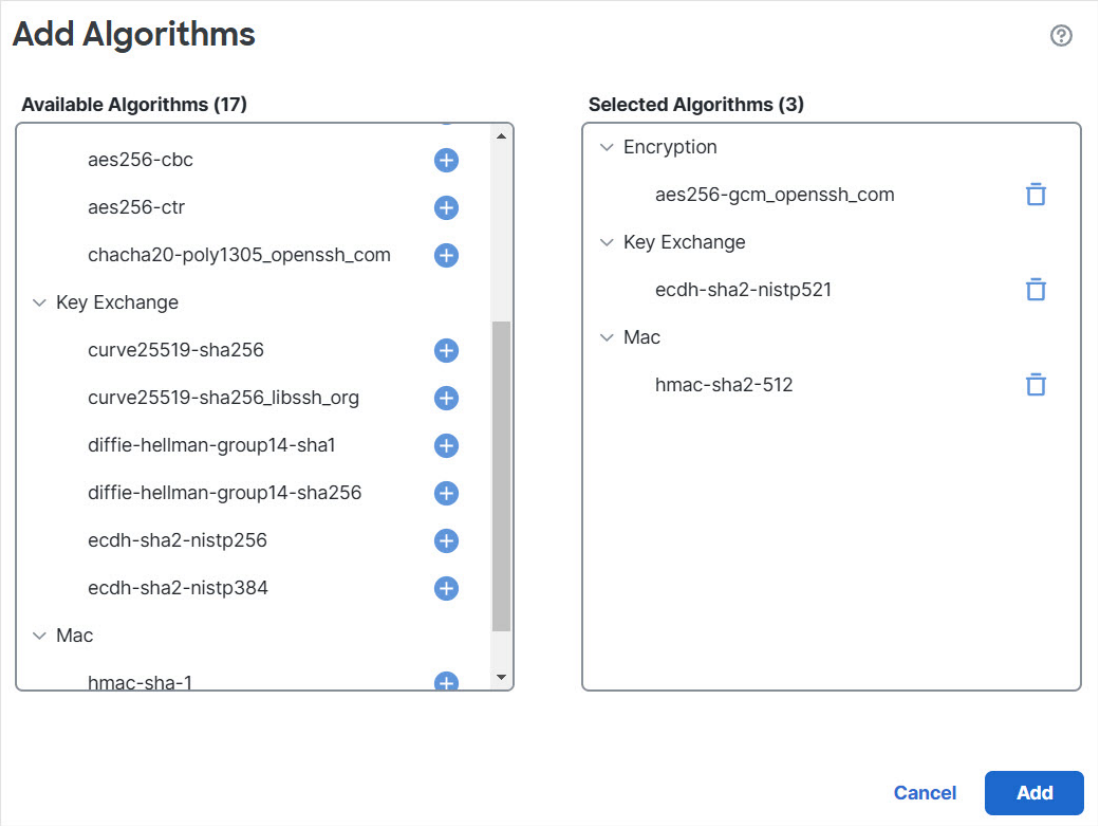

a) Select the **Encryption** algorithms.
b) Select the **Key Exchange** algorithms.

   The key exchange provides a shared secret that cannot be determined by either party alone. The key exchange is combined with a signature and the host key to provide host authentication. This key-exchange method provides explicit server authentication.

c) Select the **Mac** integrity algorithms.

Step 5 For **Host Key**, enter the modulus size for the RSA key pairs.

The modulus value (in bits) is in multiples of 8 from 1024 to 2048. The larger the key modulus size you specify, the longer it takes to generate an RSA key pair. We recommend a value of 2048.

Step 6 For the server **Volume Rekey Limit**, set the amount of traffic in KB allowed over the connection before FXOS disconnects from the session.

Step 7 For the server **Time Rekey Limit**, set the minutes for how long an SSH session can be idle before FXOS disconnects the session.

Step 8 For the **SSH Client**, configure the following settings.

Figure 179: SSH

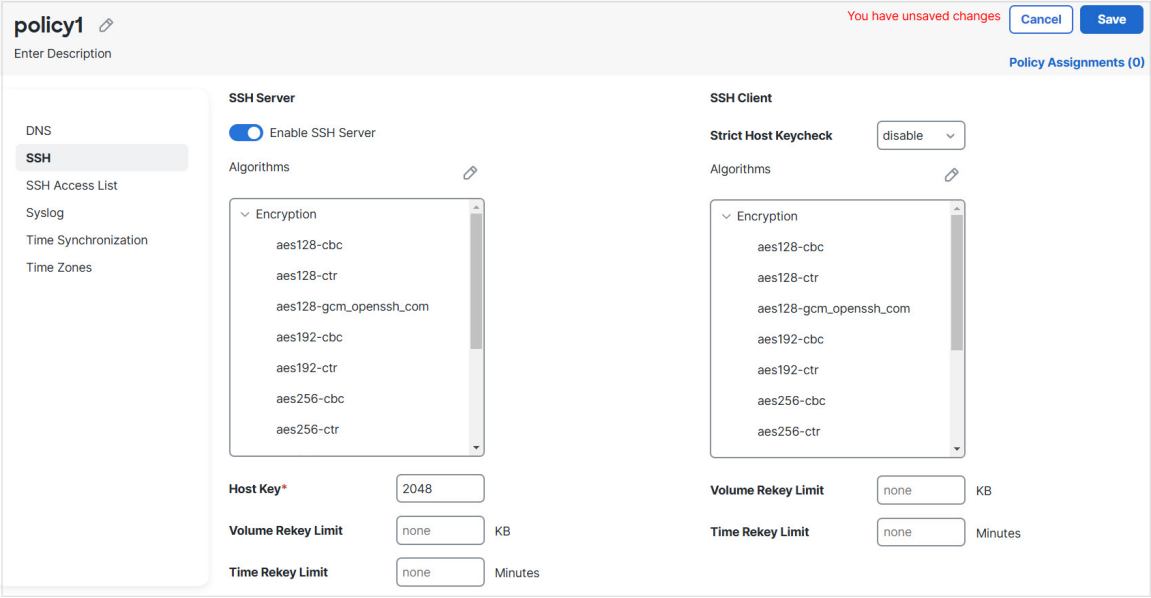

- **Strict Host Keycheck**—Choose **enable**, **disable**, or **prompt** to control SSH host key checking.
  - **enable**—The connection is rejected if the host key is not already in the FXOS known hosts file. You must manually add hosts at the FXOS CLI using the **enter ssh-host** command in the system/services scope.
  - **prompt**—You are prompted to accept or reject the host key if it is not already stored on the chassis.
  - **disable**—(The default) The chassis accepts the host key automatically if it was not stored before.
- **Algorithms**—Click **Edit** (). and select the **Encryption**, **Key Exchange**, and **Mac** algorithms.
- **Volume Rekey Limit**—Set the amount of traffic in KB allowed over the connection before FXOS disconnects from the session.
- **Time Rekey Limit**—Set the minutes for how long an SSH session can be idle before FXOS disconnects the session.

Step 9 Choose **SSH Access List**. You need to allow access to IP addresses or networks before you can use SSH.

You can add up to 25 access lists.

Figure 180: SSH Access List

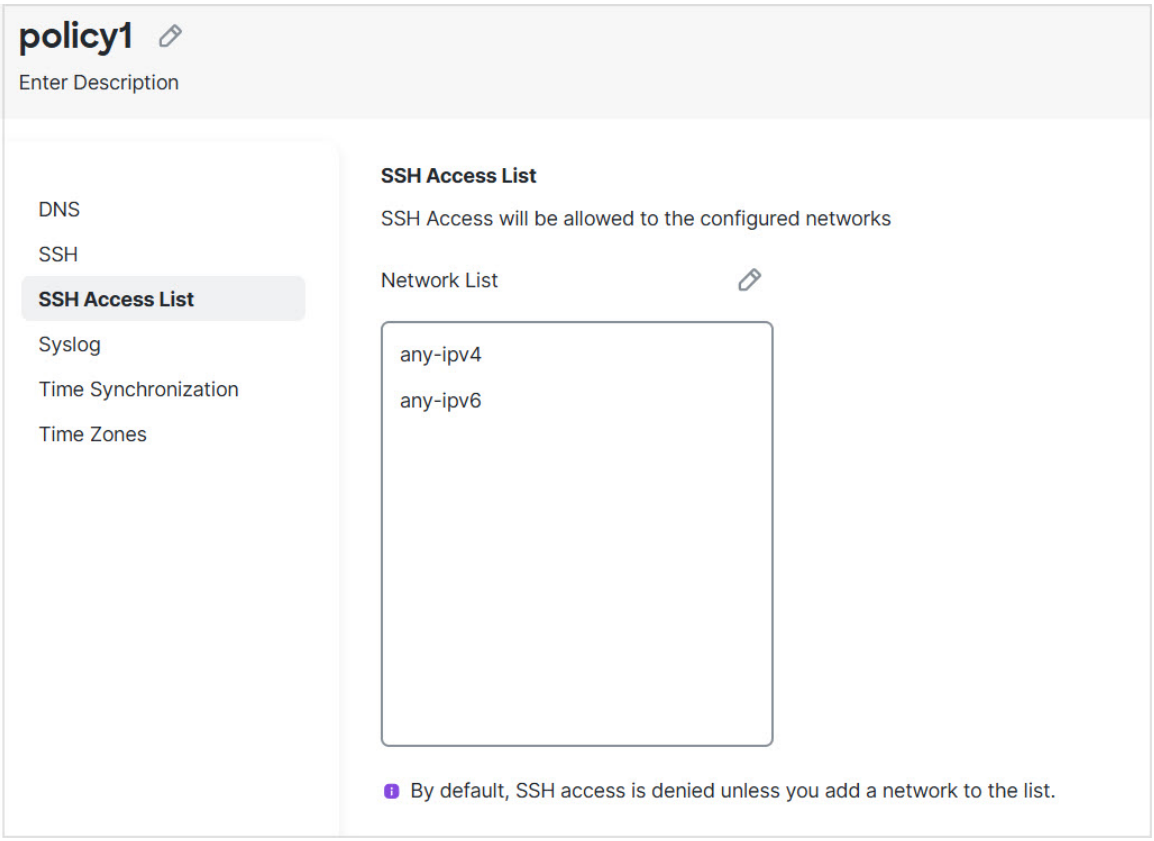

Step 10 Click **Edit** (✎) to add network objects and click **Save**. You can also manually enter IP addresses.

Figure 181: Network Objects

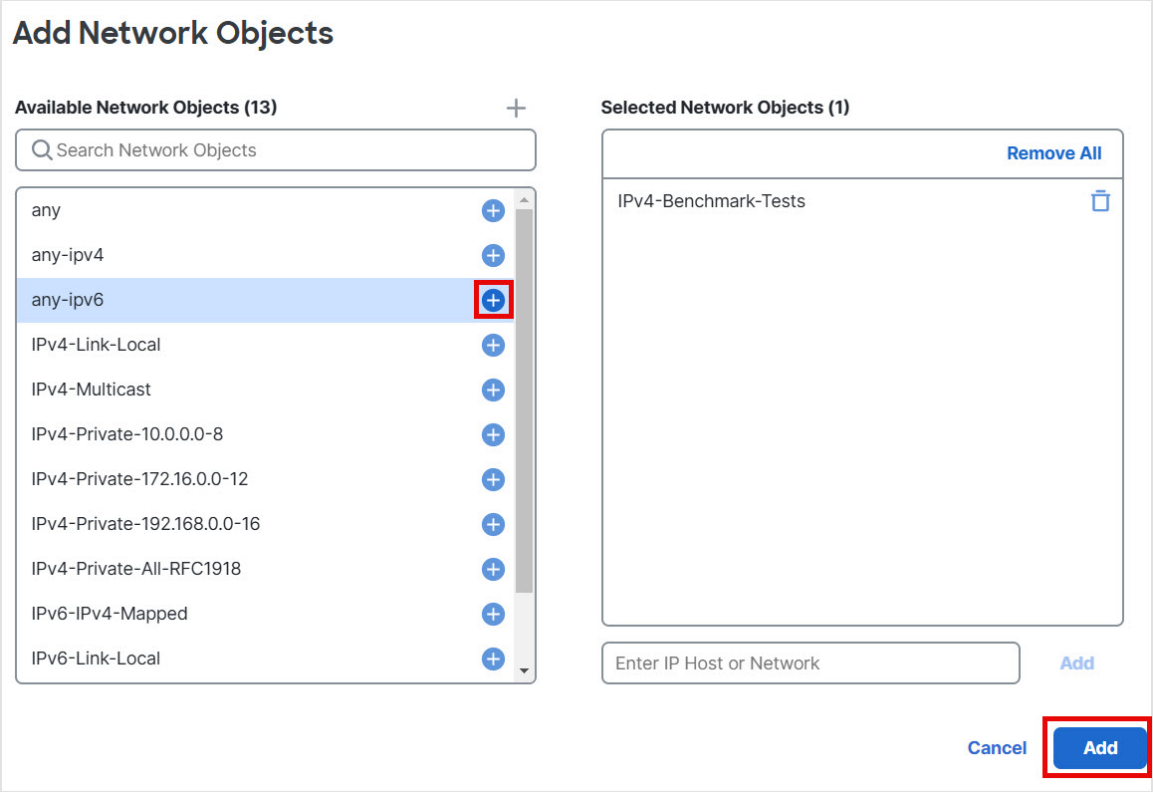

Step 11 Click **Save** to save all policy changes.

Configure Syslog

You can enable syslogs from the chassis. These syslogs come from the chassis' FXOS operating system.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the chassis policy.

Step 2 Choose **Syslog**.

Step 3 Click the **Local Destinations** tab, and complete the following fields.

Figure 182: Syslog Local Destinations

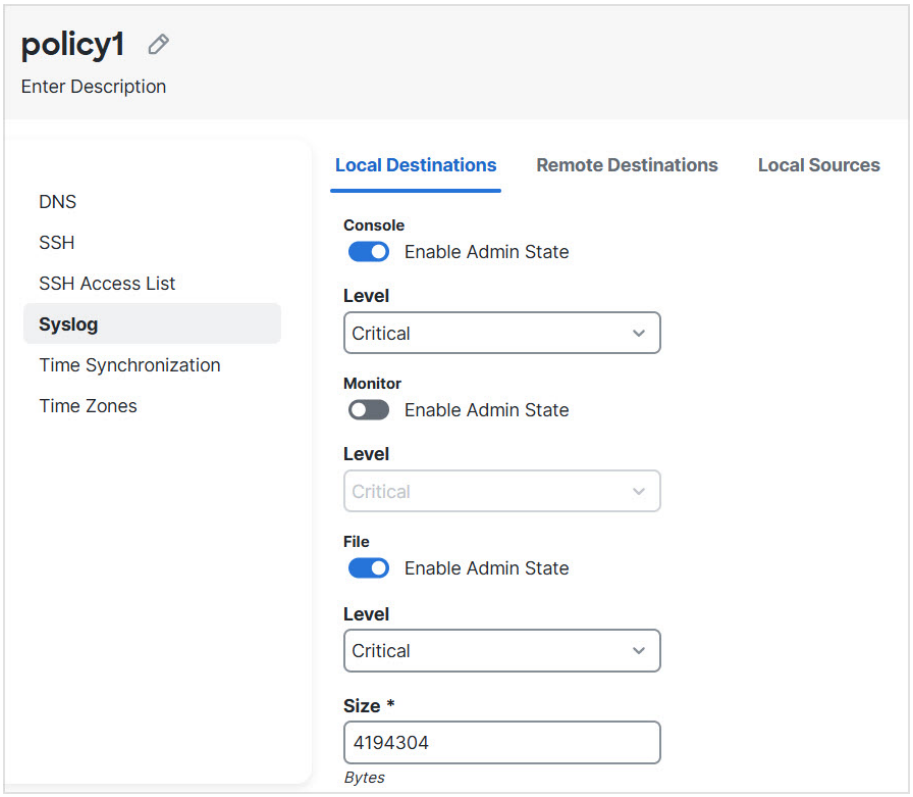

Name	Description
Console Section	
Admin State field	Whether the chassis displays syslog messages on the console. Check the **Enable** check box if you want to have syslog messages displayed on the console as well as added to the log. If the **Enable** check box is unchecked, syslog messages are added to the log but are not displayed on the console.
Level field	If you checked the **Enable** check box for **Console - Admin State**, select the lowest message level that you want displayed on the console. The chassis displays that level and above on the console. This can be one of the following: • **Emergencies** • **Alerts** • **Critical**
Monitor Section	

Name	Description
Admin State field	Whether the chassis displays syslog messages on the monitor. Check the **Enable** check box if you want to have syslog messages displayed on the monitor as well as added to the log. If the **Enable** check box is unchecked, syslog messages are added to the log but are not displayed on the monitor.
Level drop-down list	If you checked the **Enable** check box for **Monitor - Admin State**, select the lowest message level that you want displayed on the monitor. The system displays that level and above on the monitor. This can be one of the following: • **Emergencies** • **Alerts** • **Critical** • **Errors** • **Warnings** • **Notifications** • **Information** • **Debugging**

Step 4 On the **Remote Destinations** tab, complete the following fields for up to three external logs that can store messages generated by the chassis:

Figure 183: Syslog Remote Destinations

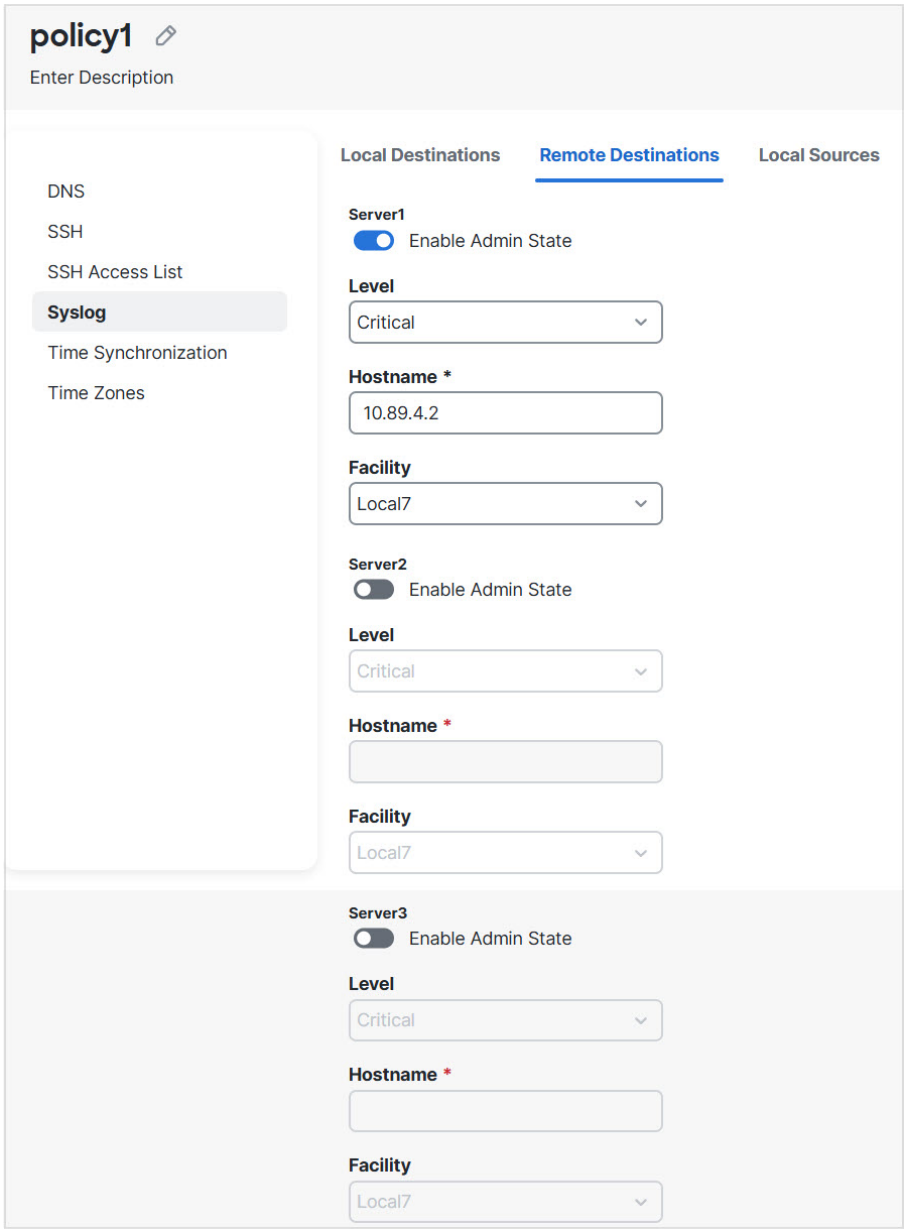

By sending syslog messages to a remote destination, you can archive messages according to the available disk space on the external syslog server, and manipulate logging data after it is saved. For example, you could specify actions to be executed when certain types of syslog messages are logged, extract data from the log and save the records to another file for reporting, or track statistics using a site-specific script.

Name	Description
Admin State field	Check the **Enable** check box if you want to have syslog messages stored in a remote log file.

Name	Description
Level drop-down list	Select the lowest message level that you want the system to store. The system stores that level and above in the remote file. This can be one of the following: • **Emergencies** • **Alerts** • **Critical** • **Errors** • **Warnings** • **Notifications** • **Information** • **Debugging**
Hostname/IP Address field	The hostname or IP address on which the remote log file resides. **Note** You must configure a DNS server if you use a hostname rather than an IP address.
Facility drop-down list	Choose a system log facility for syslog servers to use as a basis to file messages. This can be one of the following: • **Local0** • **Local1** • **Local2** • **Local3** • **Local4** • **Local5** • **Local6** • **Local7**

Step 5 Click the **Local Sources** tab, and complete the following fields.

Figure 184: Syslog Local Sources

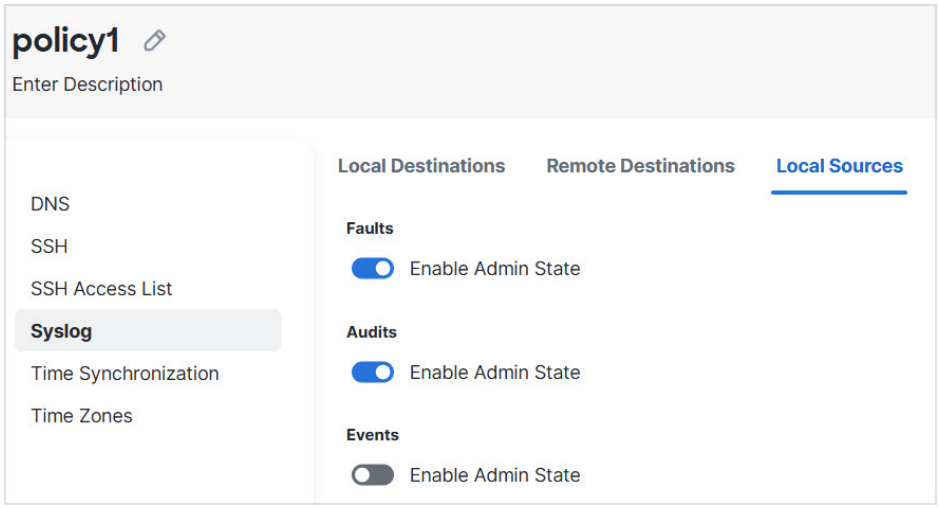

Name	Description
Faults > **Enable Admin State**	Enable system fault logging.
Audits > **Enable Admin State**	Enable audit logging.
Events > **Enable Admin State**	Enable system event logging.

Step 6 Click **Save** to save all policy changes.

Configure Time Synchronization

NTP is used to implement a hierarchical system of servers that provide a precisely synchronized time among network systems. This kind of accuracy is required for time-sensitive operations, such as validating CRLs, which include a precise time stamp. You can configure up to four NTP servers.

Note

- FXOS uses NTP version 3.
- If the stratum value of an external NTP server is 13 or greater, the application instance cannot sync to the NTP server on the FXOS chassis. Each time a NTP client syncs to a NTP server, the stratum value increases by one.

  If you have set up your own NTP server, you can find its stratum value in the /etc/ntp.conf file on the server. If the NTP server has stratum value of 13 or greater you can either change the stratum value in the ntp.conf file and restart the server, or use a different NTP server (for example: pool.ntp.org).

Before you begin

If you use a hostname for the NTP server, you must configure a DNS server. See Configure DNS, on page 396.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the chassis policy.

Step 2 Choose **Time Synchronization**.

Figure 185: Time Synchronization

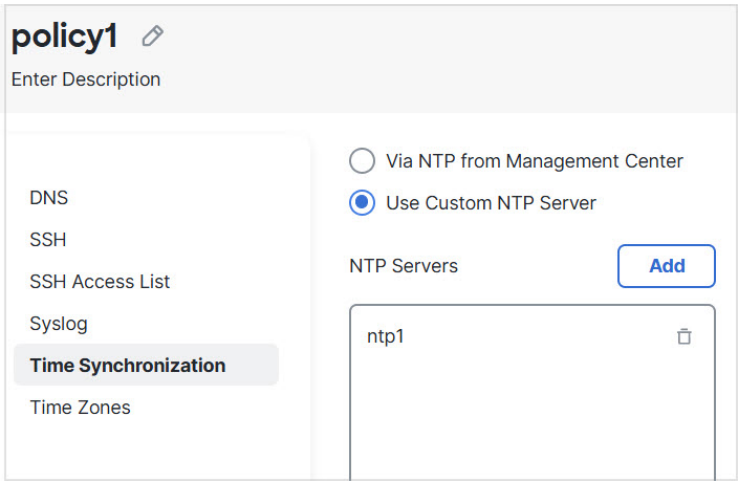

Step 3 If you want to obtain the time from the Firewall Management Center, click **Via NTP from Management Center**.

This option ensures both the chassis and the Firewall Management Center have the same time.

Step 4 To use an external NTP server, click **Use Custom NTP Server**.

a) Click **Add** to add a server.

Figure 186: Add NTP Server

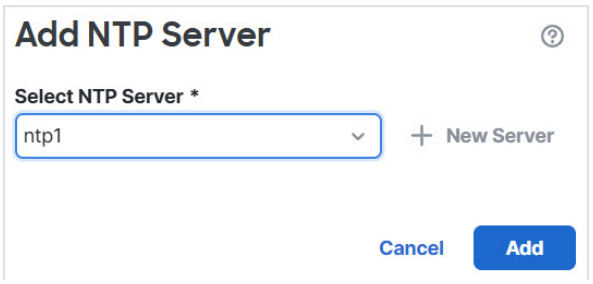

b) Choose any already-defined servers from the drop-down menu and click **Add**, or click **New Server** to add a new server.

Figure 187: Add New NTP Server

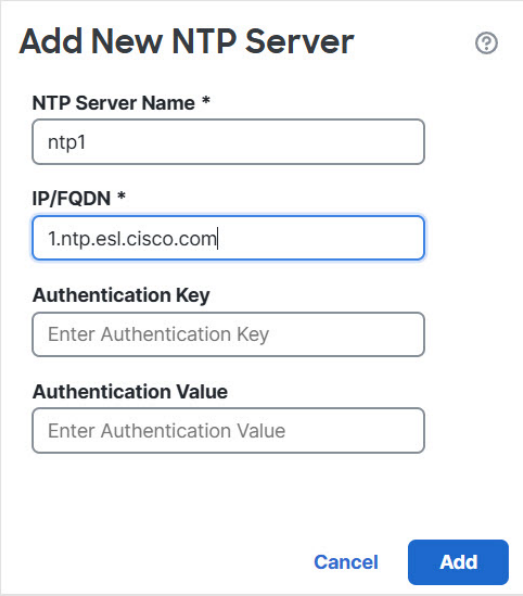

c) For a new server, enter the following fields, and click **Add**.

- **NTP Server Name**—A name to identify this server.
- **IP/FQDN**—The IP address or hostname of the server.
- **Authentication Key** and **Authentication Value**—Obtain the key ID and value from the NTP server. For example, to generate the SHA1 key on NTP server Version 4.2.8p8 or later with OpenSSL installed, enter the **ntp-keygen -M** command, and then view the key ID and value in the ntp.keys file. The key is used to tell both the client and server which value to use when computing the message digest.

  Only SHA1 is supported for NTP server authentication.

Step 5 Click **Save** to save all policy changes.

Configure Time Zones

Set the time zone for the chassis.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the chassis policy.

Step 2 Choose **Time Zones**.

Figure 188: Time Zones

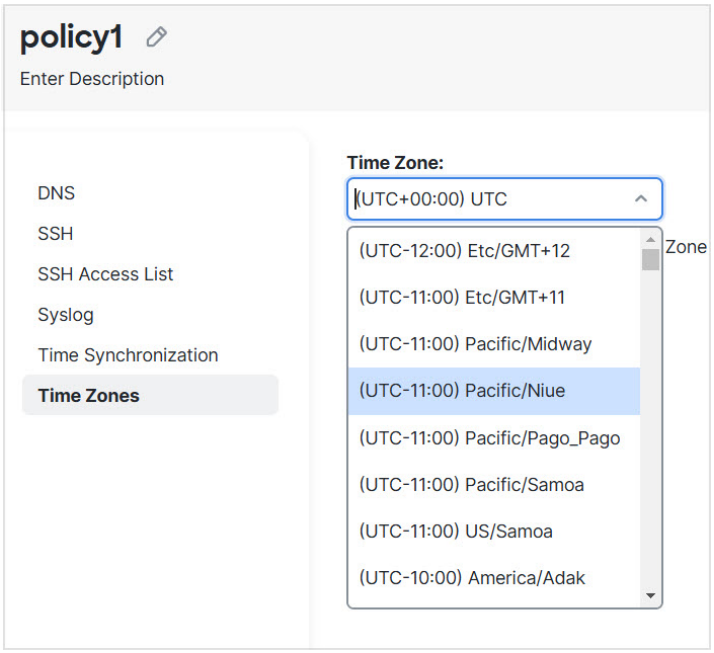

Step 3 Choose your **Time Zone** from the drop-down menu.

Step 4 Click **Save** to save all policy changes.

Manage Multi-Instance Mode

This section describes less common tasks, including changing settings at the FXOS CLI or changing interfaces assigned to the chassis.

Enable Multi-Instance Mode at the CLI

If you want to pre-configure your devices for multi-instance mode before you add them to the Firewall Management Center, you can perform this procedure. To use the Firewall Management Center to convert to multi-instance mode, see Convert a Device to Multi-Instance Mode, on page 372.

You need to connect to the Firewall Threat Defense CLI at the console port to enable multi-instance mode. After you configure the mode, you can add it to the Firewall Management Center.

Note Although you can connect to SSH on the management port, we recommend using the console port to avoid multiple disconnections. This procedure covers the console port.

Procedure

Step 1 Connect to the chassis console port.

The console port connects to the FXOS CLI.

Step 2 Log in with the username **admin** and the password **Admin123**.

The first time you log in to FXOS, you are prompted to change the password.

Note
If the password was already changed, and you do not know it, you must reimage the device to reset the password to the default. See the FXOS troubleshooting guide for the reimage procedure.

Example:

```
firepower login: admin
Password: Admin123
Successful login attempts for user 'admin' : 1

[...]

Hello admin. You must change your password.
Enter new password: ********
Confirm new password: ********
Your password was updated successfully.

[...]

firepower#
```

Step 3 Check your current mode, Native or Container. If the mode is Native, you can continue with this procedure to convert to multi-instance (Container) mode.

show system detail

Example:

```
firepower # show system detail

Systems:
    Name: firepower
    Mode: Stand Alone
    System IP Address: 172.16.0.50
    System IPv6 Address: ::
    System Owner:
    System Site:
    Deploy Mode: Native
    Description for System:
firepower #
```

Step 4 Connect to the Firewall Threat Defense CLI.

connect ftd

Example:

```
firepower# connect ftd
>
```

Step 5 The first time you log in to the Firewall Threat Defense, you are prompted to accept the General Terms. You are then presented with the CLI setup script.

The setup script lets you set the Management interface IP address and other settings. However, when you convert to multi-instance mode, the only settings that are retained are the following.

- Admin password (that you set at initial login)
- DNS servers
- Search domains

You will reset the Management IP address and gateway as part of the multi-instance mode command. After you convert to multi-instance mode, you can change Management settings at the FXOS CLI. See Change Chassis Management Settings at the FXOS CLI, on page 415.

Step 6 Enable multi-instance mode, set the chassis management interface settings, and identify the Firewall Management Center. You can use IPv4 and/or IPv6 static addressing; DHCP is not supported. After you enter the command, you will be prompted to erase the configuration and reboot. Enter **ERASE** (all caps). The system reboots and, as part of changing the mode, erases the configuration with the exception of the Management network settings you set in the command and the admin password. The chassis hostname is set to "firepower-*model*."

IPv4:

configure multi-instance network ipv4 *ip_address network_mask gateway_ip_address* **manager** *manager_name* {*hostname* | *ipv4_address* | **DONTRESOLVE**} *registration_key nat_id*

IPv6:

configure multi-instance network ipv6 *ipv6_address prefix_length gateway_ip_address* **manager** *manager_name* {*hostname* | *ipv6_address* | **DONTRESOLVE**} *registration_key nat_id*

See the following **manager** components:

- {*hostname* | *ipv4_address* | **DONTRESOLVE**}—Specifies either the FQDN or IP address of the Firewall Management Center. At least one of the devices, either the Firewall Management Center or the chassis, must have a reachable IP address to establish the two-way, SSL-encrypted communication channel between the two devices. If you don't specify a manager hostname or IP address in this command, then enter **DONTRESOLVE**; in this case, the chassis must have a reachable IP address or hostname, and you must specify the *nat_id*.
- *registration_key*—Enter a one-time registration key of your choice that you will also specify on the Firewall Management Center when you register the chassis. The registration key must not exceed 37 characters. Valid characters include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-).
- *nat_id*—Specifies a unique, one-time string of your choice that you will also specify on the Firewall Management Center when you register the chassis when one side does not specify a reachable IP address or hostname. It is required if you do not specify a manager address or hostname, however, we recommend that you always set the NAT ID even when you specify a hostname or IP address. The NAT ID must not exceed 37 characters. Valid characters include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-). This ID cannot be used for any other devices registering to the Firewall Management Center.

To change the mode back to appliance mode, you must use the FXOS CLI and enter **scope system** and then **set deploymode native**. See Change Chassis Management Settings at the FXOS CLI, on page 415.

Example:

```
> configure multi-instance network ipv4 172.16.0.104 255.255.255.0 172.16.0.1 manager
fmc1 172.16.0.103 impala67 winchester1
WARNING: This command will discard any FTD configuration (except admin's credentials).
Make sure you backup your content. All previous content will be lost. System is going to
be re-initialized.
Type ERASE to confirm:ERASE
Exit...
>
```

Step 7 Add the chassis to the Firewall Management Center. See Add a Chassis, on page 53.

Change Interfaces Assigned to an Instance

You can allocate or unallocate an interface on the instance. Adding a new interface, or deleting an unused interface, has minimal impact on the instance configuration. You can also edit the membership of an allocated EtherChannel without affecting the instance. However, deleting an interface that is used in your security policy will impact the configuration.

Interfaces can be referenced directly in many places in the instance configuration, including access rules, NAT, SSL, identity rules, VPN, DHCP server, and so on. Deleting an interface will delete any configuration associated with that interface.

Policies that refer to security zones are not affected.

Note For high availability, you need to make the same interface changes for the other unit. Otherwise, high availability might not operate correctly.

Before you begin

- Configure your interfaces according to Configure Instances, on page 372.
- If you want to add an already-allocated interface to an EtherChannel, you need to unallocate the interface from the instance first, then add the interface to the EtherChannel. For a new EtherChannel, you can then allocate the EtherChannel to the instance.

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column or click **Edit** ( ).

Figure 189: Manage Chassis

The **Chassis Manager** page opens for the chassis to the **Summary** page.

Step 2 Click **Instances**, and click **Edit** (✎) next to the instance for which you want to change interfaces.

Figure 190: Instances

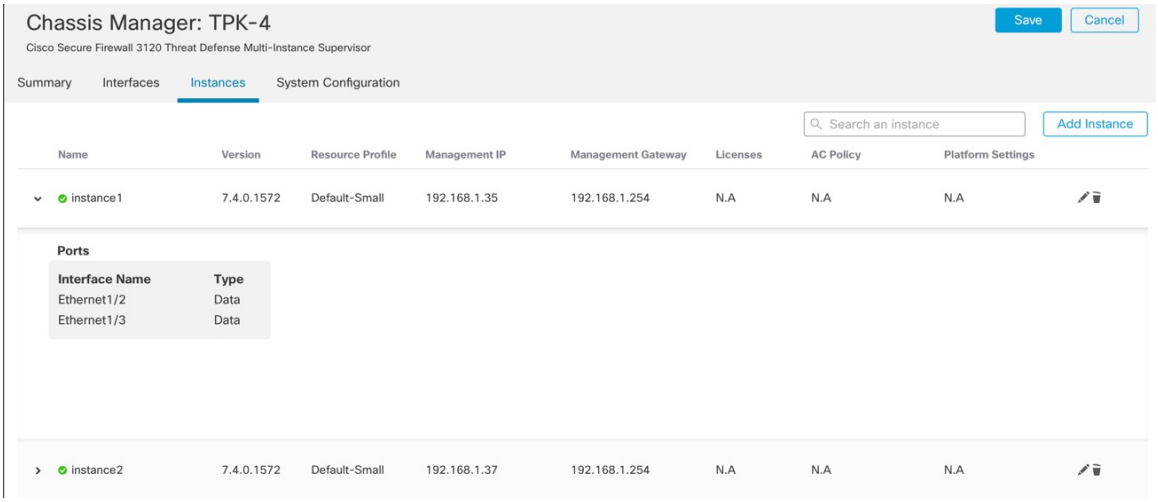

Step 3 Click **Next** until you get to the **Interface Assignment** screen.

Figure 191: Interface Assignment

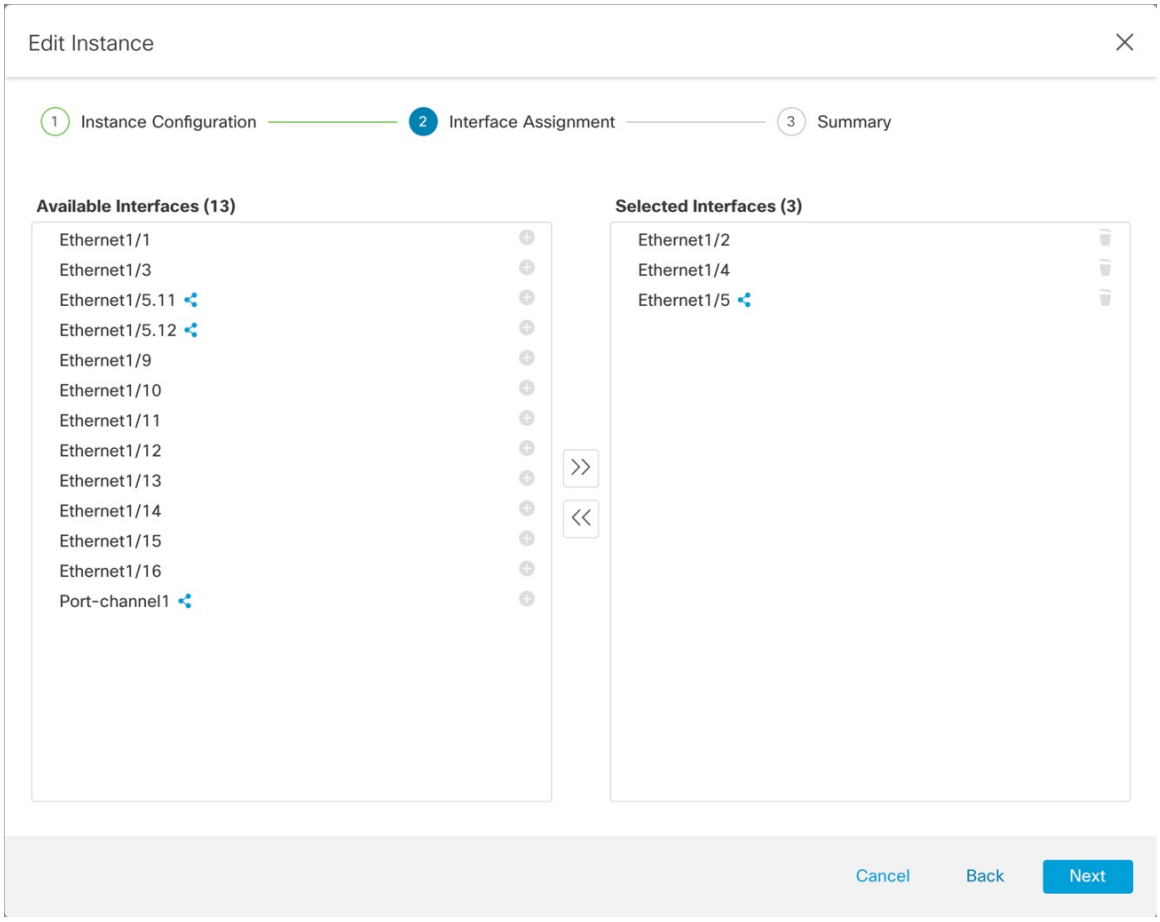

Shared interfaces show the sharing icon ( ).

Step 4 Make your interface changes, and then click **Next**.

Step 5 Click **Save** on the **Summary** screen.

Step 6 For high availability, you need to make the same interface changes for the other unit. Otherwise, high availability might not operate correctly.

Change Chassis Management Settings at the FXOS CLI

If you want to change the chassis management interface IP address and gateway, change the Firewall Management Center to a new manager, change the admin password, or disable multi-instance mode, you can do so from the FXOS CLI.

Procedure

Step 1 Connect to the chassis console port.

The console port connects to the FXOS CLI.

Note
We recommend using the console port. You can also connect using SSH to the management interface, if configured in the chassis platform settings in the Firewall Management Center; however, if you change the management IP address, you will be disconnected.

Step 2 Log in with the username **admin** and the password you set during initial setup.

Step 3 Change the Management IP address. You can use a static IPv4 and/or IPv6 address.

IPv4:

scope fabric-interconnect

set out-of-band static ip *ip_address* **netmask** *network_mask* **gw** *gateway_ip_address*

IPv6:

scope fabric-interconnect

scope ipv6-config

set out-of-band static ipv6 *ipv6_address* **ipv6-prefix** *prefix_length* **ipv6-gw** *gateway_address*

Example:

IPv4:

```
firepower-3110# scope fabric-interconnect
firepower-3110 /fabric-interconnect # set out-of-band static ip 10.5.23.8 netmask
255.255.255.0
gw 10.5.23.1
```

IPv6:

```
firepower-3110# scope fabric-interconnect
firepower-3110 / fabric-interconnect # scope ipv6-config
firepower-3110 / fabric-interconnect /ipv6-config # set out-of-band static ipv6 2001:DB8::34
ipv6-prefix 64 ipv6-gw 2001:DB8::1
```

Step 4 Change the Firewall Management Center.

You should first unregister the chassis from the current Firewall Management Center.

enter device-manager *manager_name* [**hostname** {*hostname* | *ipv4_address* | *ipv6_address*}] [**nat-id** *nat_id*]

You are prompted for the registration key.

You can enter this command from any scope.

- **hostname** {*hostname* | *ipv4_address* | *ipv6_address*}—Specifies either the FQDN or IP address of the Firewall Management Center. At least one of the devices, either the Firewall Management Center or the chassis, must have a reachable IP address to establish the two-way, TLS-1.3-encrypted communication channel between the two devices. If you do not specify a **hostname**, then the chassis must have a reachable IP address or hostname and you must specify the **nat-id**.
- **nat-id** *nat_id*—Specifies a unique, one-time string of your choice that you will also specify on the Firewall Management Center when you register the chassis when one side does not specify a reachable IP address

or hostname. It is required if you do not specify a **hostname**, however we recommend that you always set the NAT ID even when you specify a hostname or IP address. The NAT ID must not exceed 37 characters. Valid characters include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-). This ID cannot be used for any other devices registering to the Firewall Management Center.

- **Registration Key:** *reg_key*—You will be prompted for a one-time registration key of your choice that you will also specify on the Firewall Management Center when you register the chassis. The registration key must not exceed 37 characters. Valid characters include alphanumerical characters (A–Z, a–z, 0–9) and the hyphen (-).

Example:

```
firepower-3110# enter device-manager boulder_fmc hostname 10.89.5.35 nat-id 93002
(Valid registration key characters: [a-z],[A-Z],[0-9],[-]. Length: [2-36])
Registration Key: Impala67
```

Step 5 Change the admin password.

scope security

set password

Enter a password: *password*

Confirm the password: *password*

Example:

```
firepower-3110# scope security
firepower-3110 /security # set password
Enter new password: Sw@nsong67
Confirm new password: Sw@nsong67
firepower-3110 /security #
```

Step 6 Disable multi-instance mode and set the system back to appliance mode.

scope system

set deploymode native

You are prompted to reboot.

Example:

```
firepower-3110# scope system
firepower-3110 /system # set deploymode native
All configuration and bootable images will be lost and system will reboot.
If there was out of band upgrade, it might reboot with the base version and
need to re-image to get the expected running version.
Do you still want to change deploy mode? (yes/no):yes
firepower-3110 /system #
```

To change the mode back to multi-instance mode, enter **set deploymode container**. You can check the current mode using the **show system detail** command.

Monitoring Multi-Instance Mode

This section helps you troubleshoot and diagnose your multi-instance mode chassis and instances.

Monitoring Multi-Instance Setup

show system detail

This FXOS command shows the current mode, Native or Container. If the mode is Native (also known as appliance mode), you can convert to multi-instance (Container) mode. Note that the prompt/name in multi-instance mode is generic "firepower-<model>" while the prompt in appliance mode is the hostname you set for the Firewall Threat Defense (by default "firepower")

```
firepower # show system detail

Systems:
    Name: firepower
    Mode: Stand Alone
    System IP Address: 172.16.0.50
    System IPv6 Address: ::
    System Owner:
    System Site:
    Deploy Mode: Native
    Description for System:
firepower #
```

scope system > show

This FXOS command shows the current mode in a table format. Note that the prompt/name in multi-instance mode is generic "firepower-<model>" while the prompt in appliance mode is the hostname you set for the Firewall Threat Defense.

```
firepower-3110# scope system
firepower-3110 /system # show

Systems:
    Name       Mode        Deploy Mode System IP Address System IPv6 Address
    ---------- ----------- ----------- ----------------- -------------------
    firepower-3110
               Stand Alone Container   10.89.5.42        ::


3110-1# scope system
3110-1 /system # show

Systems:
    Name       Mode        Deploy Mode System IP Address System IPv6 Address
    ---------- ----------- ----------- ----------------- -------------------
    3110-1     Stand Alone Native      10.89.5.41        ::
3110-1 /system #
```

Monitoring Instance Interfaces

show portmanager switch forward-rules hardware mac-filter

This command shows the internal switch-forwarding rule for two instances with a dedicated physical interface assigned to each instance. Ethernet 1/2 is assigned to ftd1 and Ethernet 1/1 is assigned to ftd2.

ECMP group 1540 is assigned to ftd1 and ECMP group 1541 is assigned to ftd2.

```
secfw-3140(local-mgmt)# show portmanager switch forward-rules hardware mac-filter
      VLAN   SRC_PORT   PC_ID   SRC_ID   DST_PORT     PKT_CNT      DMAC
1        0         17       0       17         19       29164  0:0:0:0:0:0
2        0         19       0       19         17       67588  0:0:0:0:0:0
3        0          1       0      101       1541           0  a2:5b:83:0:0:15
4        0          1       0      101       1541        8181  ff:ff:ff:ff:ff:ff
5        0          2       0      102       1540           0  a2:5b:83:0:0:18
6        0          2       0      102       1540         431  ff:ff:ff:ff:ff:ff
7        0         17       0        0          0       11133  0:0:0:0:0:0
8        0         17       0        0          0           0  0:0:0:0:0:0
```

This command shows the internal switch-forwarding rule for two instances with shared physical interfaces assigned to two instances. Ethernet 1/1 is shared between ftd1 and ftd2.

ECMP group 1540 is assigned to ftd1 and ECMP group 1541 is assigned to ftd2.

MCAST group 4096 is used for replicating broadcast traffic between ftd1 and ftd2.

```
firepower-3140(local-mgmt)# show portmanager switch forward-rules hardware mac-filter
      VLAN   SRC_PORT   PC_ID   SRC_ID   DST_PORT     PKT_CNT      DMAC
1        0         17       0       17         19        2268  0:0:0:0:0:0
2        0         19       0       19         17        4844  0:0:0:0:0:0
3        0          1       0      101       1541           0  a2:5b:83:0:0:9
4        0          1       0      101       4096         546  ff:ff:ff:ff:ff:ff
5        0          1       0      101       1540           0  a2:5b:83:0:0:c
6        0         17       0        0          0        1263  0:0:0:0:0:0
7        0         17       0        0          0           0  0:0:0:0:0:0
```

This command shows the internal switch-forwarding rule for two instances with shared subinterfaces assigned to both instances. Ethernet 1/1.2452 is shared between ftd1 and ftd2.

ECMP group 1540 is assigned to ftd1 and ECMP group 1541 is assigned to ftd2.

MCAST group 4097 is used for replicating broadcast traffic between ftd1 and ftd2.

```
firepower-3140(local-mgmt)# show portmanager switch forward-rules hardware mac-filter
      VLAN   SRC_PORT   PC_ID   SRC_ID   DST_PORT     PKT_CNT      DMAC
1        0         17       0       17         19       21305  0:0:0:0:0:0
2        0         19       0       19         17       50976  0:0:0:0:0:0
3     2452          1       0      101       1541         430  a2:5b:83:0:0:f
4     2452          1       0      101       4097           0  ff:ff:ff:ff:ff:ff
5     2452          1       0      101       1540           0  a2:5b:83:0:0:12
6        0         17       0        0          0       11038  0:0:0:0:0:0
7        0         17       0        0          0           0  0:0:0:0:0:0
```

show portmanager switch ecmp-groups detail

Use this command to list each Instance Ecmp-Vport-Physical port mapping detail.

Note Physical-Port 18 is the backplane uplink interface between the internal switch and the instance.

```
firepower-3140(local-mgmt)# show portmanager switch ecmp-groups detail
       ECMP-GROUP  VPORT      PHYSICAL-PORT
1      1536        256        18
2      1537        257        18
3      1538        258        18
4      1539        259        18
5      1540        260        18
6      1541        261        18
7      1542        262        18
8      1543        263        18
9      1544        264        18
10     1545        265        18
```

show portmanager switch mcast-groups detail

Use this command to list MCAST group membership details.

```
firepower-3140(local-mgmt)# show portmanager switch mcast-groups detail
       MCAST-GROUP
1      4096
       Member-ports
       Ethernet 1/1
       ECMP-ID 1541
       ECMP-ID 1540
```

show portmanager counters mcast-group

Use this command to check the MCAST group packet counter.

```
firepower-3140(local-mgmt)# show portmanager counters mcast-group 4096
PKT_CNT: 8106
```

show portmanager counters ecmp

Use this command to check the ECMP group packet counter.

```
firepower-3140(local-mgmt)# show portmanager counters ecmp 1541
PKT_CNT: 430
```

History for Multi-Instance Mode

Table 26:

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Multi-instance mode conversion in the Firewall Management Center	7.6.0	7.6.0	You can now register an application-mode device to the Firewall Management Center and then convert it to multi-instance mode without having to use the CLI. New/modified screens: • **Devices** > **Device Management**, then for a device, click **More** (⋮) > **Convert to Multi-Instance** • **Devices** > **Device Management**, then select multiple devices and choose **Select Bulk Action** > **Convert to Multi-Instance**
Multi-instance mode for the Secure Firewall 4200	7.6.0	7.6.0	Multi-instance mode is now supported for the Secure Firewall 4200.
Multi-instance mode for the Secure Firewall 3100.	7.4.1	7.4.1	You can deploy the Secure Firewall 3100 as a single device (*appliance mode*) or as multiple container instances (*multi-instance mode*). In multi-instance mode, you can deploy multiple container instances on a single chassis that act as completely independent devices. Note that in multi-instance mode, you upgrade the operating system and the firmware (*chassis upgrade*) separately from the container instances (*Firewall Threat Defense upgrade*). New/modified screens: • **Devices** > **Device Management** > **Add** > **Chassis** • **Devices** > **Device Management** > **Device** > **Chassis Manager** • **Devices** > **Platform Settings** > **New Policy** > **Chassis Platform Settings** • **Devices** > **Chassis Upgrade** New/modified Firewall Threat Defense CLI commands: **configure multi-instance network ipv4**, **configure multi-instance network ipv6** New/modified FXOS CLI commands: **create device-manager**, **set deploymode** Platform restrictions: Not supported on the Secure Firewall 3105.

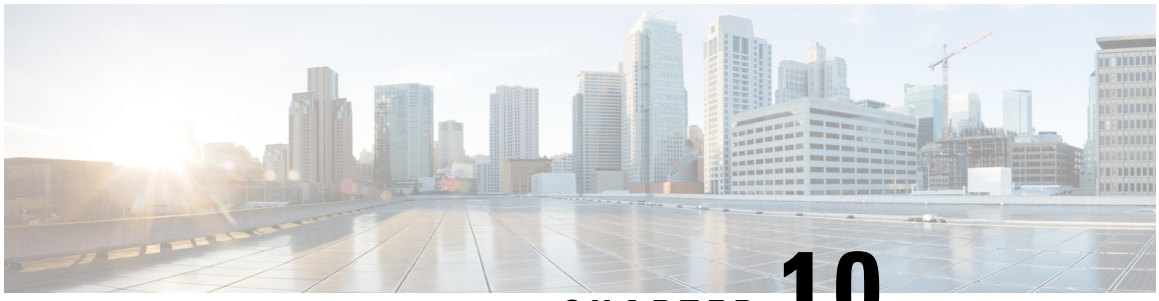

CHAPTER 10

High Availability

The following topics describe how to configure Active/Standby failover to accomplish high availability of the Firewall Threat Defense.

- About Secure Firewall Threat Defense High Availability, on page 423
- Config-Sync Optimization, on page 437
- Requirements and Prerequisites for High Availability, on page 438
- Guidelines for High Availability, on page 439
- Add a High Availability Pair, on page 441
- Configure Optional High Availability Parameters, on page 444
- Manage High Availability, on page 446
- Monitoring High Availability, on page 453
- Troubleshoot Configuration Sync Failure, on page 453
- History for High Availability, on page 454

About Secure Firewall Threat Defense High Availability

Configuring high availability, also called failover, requires two identical Firewall Threat Defense devices connected to each other through a dedicated failover link and, optionally, a state link. Firewall Threat Defense supports Active/Standby failover, where one unit is the active unit and passes traffic. The standby unit does not actively pass traffic, but synchronizes configuration and other state information from the active unit. When a failover occurs, the active unit fails over to the standby unit, which then becomes active.

The health of the active unit (hardware, interfaces, software, and environmental status) is monitored to determine if specific failover conditions are met. If those conditions are met, failover occurs.

Note High availability is not supported on Firewall Threat Defense Virtual running in the public cloud. See the Cisco Secure Firewall Threat Defense Virtual Getting Started Guide for more information about configuring the Firewall Threat Defense Virtual device for high availability.

High Availability System Requirements

This section describes the hardware, software, and license requirements for Firewall Threat Defense devices in a High Availability configuration.

Hardware Requirements

The two units in a High Availability configuration must:

- Be the same model. In addition, for container instances, they must use the same resource profile attributes.

  For the Firepower 9300, High Availability is only supported between same-type modules; but the two chassis can include mixed modules. For example, each chassis has an SM-56, SM-48, and SM-40. You can create High Availability pairs between the SM-56 modules, between the SM-48 modules, and between the SM-40 modules.

  If you change the resource profile after you add the High Availability pair to the Firewall Management Center, update the inventory for each unit on the **Devices** > **Device Management** > **Device** > **System** > **Inventory** dialog box.

  If you assign a different profile to instances in an established high-availability pair, which requires the profile to be the same on both units, you must:

  1. Break high availability.
  2. Assign the new profile to both units.
  3. Re-establish high availability.

- Have the same number and types of interfaces.

  For the Firepower 4100/9300 chassis, all interfaces must be preconfigured in FXOS identically before you enable High Availability. If you change the interfaces after you enable High Availability, make the interface changes in FXOS on the Standby unit, and then make the same changes on the Active unit.

If you are using units with different flash memory sizes in your High Availability configuration, make sure the unit with the smaller flash memory has enough space to accommodate the software image files and the configuration files. If it does not, configuration synchronization from the unit with the larger flash memory to the unit with the smaller flash memory will fail.

Software Requirements

The two units in a High Availability configuration must:

- Be in the same firewall mode (routed or transparent).
- Have the same software version.
- Be in the same domain or group on the Firewall Management Center.
- Have the same NTP configuration. See Configure NTP Time Synchronization for Threat Defense.
- Be fully deployed on the Firewall Management Center with no uncommitted changes.
- Not have DHCP or PPPoE configured in any of their interfaces.
- (Firepower 4100/9300) Have the same flow offload mode, either both enabled or both disabled.

License Requirements for Firewall Threat Defense Devices in a High Availability Pair

Both Firewall Threat Defense units in a high availability configuration must have the same licenses.

High availability configurations require two license entitlements: one for each device in the pair.

Before high availability is established, it does not matter which licenses are assigned to the secondary/standby device. During high availability configuration, the Firewall Management Center releases any unnecessary licenses assigned to the standby unit and replaces them with identical licenses assigned to the primary/active unit. For example, if the active unit has a Essentials license and a IPS license, and the standby unit has only a Essentials license, the Firewall Management Center communicates with the Smart Software Manager to obtain an available IPS license from your account for the standby unit. If your license account does not include enough purchased entitlements, your account becomes Out-of-Compliance until you purchase the correct number of licenses.

Failover and Stateful Failover Links

The failover link and the optional stateful failover link are dedicated connections between the two units. Cisco recommends to use the same interface between two devices in a failover link or a stateful failover link. For example, in a failover link, if you have used eth0 in device 1, use the same interface (eth0) in device 2 as well.

Failover Link

The two units in a failover pair constantly communicate over a failover link to determine the operating status of each unit.

Failover Link Data

The following information is communicated over the failover link:

- The unit state (active or standby)
- Hello messages (keep-alives)
- Network link status
- MAC address exchange
- Configuration replication and synchronization

Interface for the Failover Link

You can use an unused data interface (physical, or EtherChannel) as the failover link; however, you cannot specify an interface that is currently configured with a name. You also cannot use a subinterface with the exception of a subinterface defined on the chassis for multi-instance mode. The failover link interface is not configured as a normal networking interface; it exists for failover communication only. This interface can only be used for the failover link (and also for the state link).

The Firewall Threat Defense does not support sharing interfaces between user data and the failover link. You also cannot use separate subinterfaces on the same parent for the failover link and for data (multi-instance chassis subinterfaces only). If you use a chassis subinterface for the failover link, then all subinterfaces on that parent, and the parent itself, are restricted for use as failover links.

Note When using an EtherChannel as the failover or state link, you must confirm that the same EtherChannel with the same member interfaces exists on both devices before establishing high availability.

See the following guidelines for the failover link:

- Firepower 4100/9300—You cannot use the management-type interface for the failover link.
- See the following guidelines for sizing the link.

Table 27: Failover Link Size

Model	Interface Size for Combined Failover and State Link
Firepower 1010	1 Gbps
Firepower 1100	1 Gbps
Secure Firewall 1200	1 Gbps
Secure Firewall 3100	Secure Firewall 3105—1 Gbps Secure Firewall 3110—1 Gbps Secure Firewall 3120—1 Gbps Secure Firewall 3130—10 Gbps Secure Firewall 3140—10 Gbps
Firepower 4100	10 Gbps
Secure Firewall 4200	10 Gbps
Firepower 9300	10 Gbps

The alternation frequency is equal to the unit hold time.

Note If you have a large configuration and a low unit hold time, alternating between the member interfaces can prevent the secondary unit from joining/re-joining. In this case, disable one of the member interfaces until after the secondary unit joins.

For an EtherChannel used as the failover link, to prevent out-of-order packets, only one interface in the EtherChannel is used. If that interface fails, then the next interface in the EtherChannel is used. You cannot alter the EtherChannel configuration while it is in use as a failover link.

Connecting the Failover Link

Connect the failover link in one of the following two ways:

- Using a switch, with no other device on the same network segment (broadcast domain or VLAN) as the failover interfaces of the Firewall Threat Defense device.
- Using an Ethernet cable to connect the units directly, without the need for an external switch.

If you do not use a switch between the units, if the interface fails, the link is brought down on both peers. This condition may hamper troubleshooting efforts because you cannot easily determine which unit has the failed interface and caused the link to come down.

Stateful Failover Link

To use Stateful Failover, you must configure a Stateful Failover link (also known as the state link) to pass connection state information.

Shared with the Failover Link

Sharing a failover link is the best way to conserve interfaces. However, you must consider a dedicated interface for the state link and failover link, if you have a large configuration and a high traffic network.

Dedicated Interface for the Stateful Failover Link

You can use a dedicated data interface (physical or EtherChannel) for the state link. See Interface for the Failover Link, on page 425 for requirements for a dedicated state link, and Connecting the Failover Link, on page 426 for information about connecting the state link as well.

For optimum performance when using long distance failover, the latency for the state link should be less than 10 milliseconds and no more than 250 milliseconds. If latency is more than 10 milliseconds, some performance degradation occurs due to retransmission of failover messages.

Avoiding Interrupted Failover and Data Links

We recommend that failover links and data interfaces travel through different paths to decrease the chance that all interfaces fail at the same time. If the failover link is down, the failover operation is suspended until the health of the failover link is restored.

See the following connection scenarios to design a resilient failover network.

Scenario 1—Not Recommended

If a single switch or a set of switches are used to connect both failover and data interfaces between two Firewall Threat Defense devices, then when a switch or inter-switch-link is down, both Firewall Threat Defense devices become active. Therefore, the two connection methods shown in the following figures are **not** recommended.

Figure 192: Connecting with a Single Switch Not Recommended

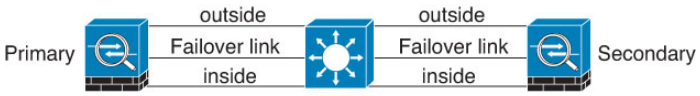

Figure 193: Connecting with a Double-Switch—Not Recommended

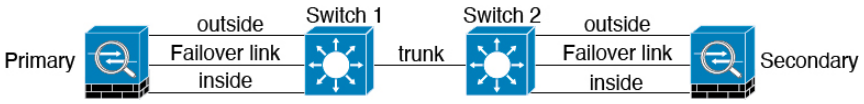

Scenario 2—Recommended

We recommend that failover links not use the same switch as the data interfaces. Instead, use a different switch or use a direct cable to connect the failover link, as shown in the following figures.

Figure 194: Connecting with a Different Switch

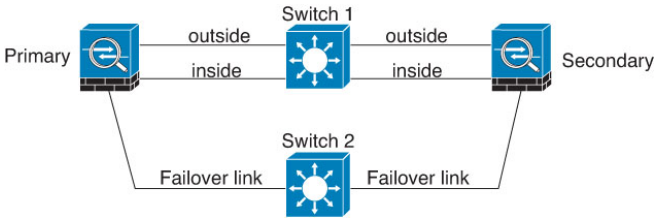

Figure 195: Connecting with a Cable

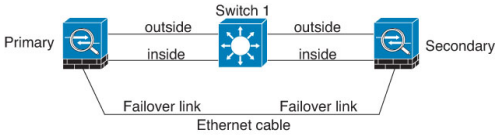

Scenario 3—Recommended

If the Firewall Threat Defense data interfaces are connected to more than one set of switches, then a failover link can be connected to one of the switches, preferably the switch on the secure (inside) side of network, as shown in the following figure.

Figure 196: Connecting with a Secure Switch

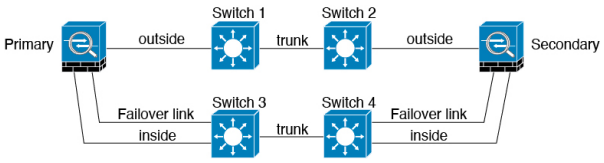

MAC Addresses and IP Addresses in High Availability

When you configure your interfaces, you can specify an active IP address and a standby IP address on the same network. Generally, when a failover occurs, the new active unit takes over the active IP addresses and MAC addresses. Because network devices see no change in the MAC to IP address pairing, no ARP entries change or time out anywhere on the network.

Note Although recommended, the standby address is not required. Without a standby IP address, the active unit cannot perform network tests to check the standby interface health; it can only track the link state. You also cannot connect to the standby unit on that interface for management purposes.

The IP address and MAC address for the state link do not change at failover.

Active/Standby IP Addresses and MAC Addresses

For Active/Standby High Availability, see the following for IP address and MAC address usage during a failover event:

1. The active unit always uses the primary unit's IP addresses and MAC addresses.

2. When the active unit fails over, the standby unit assumes the IP addresses and MAC addresses of the failed unit and begins passing traffic.
3. When the failed unit comes back online, it is now in a standby state and takes over the standby IP addresses and MAC addresses.

However, if the secondary unit boots without detecting the primary unit, then the secondary unit becomes the active unit and uses its own MAC addresses, because it does not know the primary unit MAC addresses. When the primary unit becomes available, the secondary (active) unit changes the MAC addresses to those of the primary unit, which can cause an interruption in your network traffic. Similarly, if you swap out the primary unit with new hardware, a new MAC address is used.

If you disable high availability and set the failover configurations to a disabled state, you will need to manually resume high availability, or reboot the device. It is recommended to use the command **configure high-availability resume** and resume the high availability instead of rebooting the device. If you reload the standby unit with the failover configuration disabled, the standby unit boots up as the active unit and uses the primary unit's IP addresses and MAC addresses. This leads to duplicate IP addresses and causes network traffic disruptions. Use the command **configure high-availability resume** to enable failover and restore the traffic flow.

Note If you enable failover on a standalone device, the data interfaces go down at negotiation state of failover, interrupting traffic.

Virtual MAC addresses guard against this disruption, because the active MAC addresses are known to the secondary unit at startup, and remain the same in the case of new primary unit hardware. We recommend that you configure the virtual MAC address on both the primary and secondary units to ensure that the secondary unit uses the correct MAC addresses when it is the active unit, even if it comes online before the primary unit. If you do not configure virtual MAC addresses, you might need to clear the ARP tables on connected routers to restore traffic flow. The Firewall Threat Defense device does not send gratuitous ARPs for static NAT addresses when the MAC address changes, so connected routers do not learn of the MAC address change for these addresses.

Virtual MAC Addresses

The Firewall Threat Defense device has multiple methods to configure virtual MAC addresses. We recommend using only one method. If you set the MAC address using multiple methods, the MAC address used depends on many variables, and might not be predictable.

For multi-instance capability, the FXOS chassis autogenerates only primary MAC addresses for all interfaces. You can overwrite the generated MAC address with a virtual MAC address with both the primary and secondary MAC addresses, but predefining the secondary MAC address is not essential; setting the secondary MAC address does ensure that to-the-box management traffic is not interrupted in the case of new secondary unit hardware.

MAC Address Table Update in Failover

During failover, the device that is designated as the new active device generates multicast packets for each MAC address entry in the MAC table and sends them to all the bridge group interfaces. This action prompts the upstream switches in the bridge group to update their routing tables with the new active device's interface to ensure accurate traffic forwarding.

The time taken to generate multicast packets and update the routing tables of the upstream switches depends on the number of entries in the MAC address table and the number of bridge group interfaces. Use the **show failover statistics state-switch-delay** command to display statistics related to the delays encountered during failover events.

Stateful Failover

During Stateful Failover, the active unit continually passes per-connection state information to the standby unit. After a failover occurs, the same connection information is available at the new active unit. Supported end-user applications are not required to reconnect to keep the same communication session.

Supported Features

For Stateful Failover, the following state information is passed to the standby Firewall Threat Defense device:

- NAT translation table.
- TCP and UDP connections and states, including HTTP connection states. Other types of IP protocols, and ICMP, are not parsed by the active unit, because they get established on the new active unit when a new packet arrives.
- Snort connection states, inspection results, and pin hole information, including strict TCP enforcement.
- The ARP table
- The Layer 2 bridge table (for bridge groups)
- The ISAKMP and IPsec SA table
- GTP PDP connection database
- SIP signaling sessions and pin holes.
- Static and dynamic routing tables—Stateful Failover participates in dynamic routing protocols, like OSPF and EIGRP, so routes that are learned through dynamic routing protocols on the active unit are maintained in a Routing Information Base (RIB) table on the standby unit. Upon a failover event, packets travel normally with minimal disruption to traffic because the active secondary unit initially has rules that mirror the primary unit. Immediately after failover, the re-convergence timer starts on the newly active unit. Then the epoch number for the RIB table increments. During re-convergence, OSPF and EIGRP routes become updated with a new epoch number. Once the timer is expired, stale route entries (determined by the epoch number) are removed from the table. The RIB then contains the newest routing protocol forwarding information on the newly active unit.

Note Routes are synchronized only for link-up or link-down events on an active unit. If the link goes up or down on the standby unit, dynamic routes sent from the active unit may be lost. This is normal, expected behavior.

- DHCP Server—DHCP address leases are not replicated. However, a DHCP server configured on an interface will send a ping to make sure an address is not being used before granting the address to a DHCP client, so there is no impact to the service. State information is not relevant for DHCP relay or DDNS.

- Access control policy decisions—Decisions related to traffic matching (including URL, URL category, geolocation, and so forth), intrusion detection, malware, and file type are preserved during failover. However, for connections being evaluated at the moment of failover, there are the following caveats:
  - AVC—App-ID verdicts are replicated, but not detection states. Proper synchronization occurs as long as the App-ID verdicts are complete and synchronized before failover occurs.
  - Intrusion detection state—Upon failover, once mid-flow pickup occurs, new inspections are completed, but old states are lost.
  - File malware blocking—The file disposition must become available before failover.
  - File type detection and blocking—The file type must be identified before failover. If failover occurs while the original active device is identifying the file, the file type is not synchronized. Even if your file policy blocks that file type, the new active device downloads the file.
- User identity decisions from the identity policy, including the user-to-IP address mappings gathered passively through ISE Session Directory, and active authentication through captive portal. Users who are actively authenticating at the moment of failover might be prompted to authenticate again.
- Network AMP—Cloud lookups are independent from each device, so failover does not affect this feature in general. Specifically:
  - Signature Lookup—If failover occurs in the middle of a file transmission, no file event is generated and no detection occurs.
  - File Storage—If failover occurs when the file is being stored, it is stored on the original active device. If the original active device went down while the file was being stored, the file does not get stored.
  - File Pre-classification (Local Analysis)—If failover occurs in the middle of pre-classification, detection fails.
  - File Dynamic Analysis (Connectivity to the cloud)—If failover occurs, the system might submit the file to the cloud.
  - Archive File Support—If failover occurs in the middle of an analysis, the system loses visibility into the file/archive.
  - Custom Blocking—If failover occurs, no events are generated.
- Security Intelligence decisions. However, DNS-based decisions that are in process at the moment of failover are not completed.
- RA VPN—Remote access VPN end users do not have to reauthenticate or reconnect the VPN session after a failover. However, applications operating over the VPN connection could lose packets during the failover process and not recover from the packet loss.
- From all the connections, only established ones will be replicated on the Standby device.

Unsupported Features

For Stateful Failover, the following state information is not passed to the standby Firewall Threat Defense device:

- Sessions in plaintext tunnels other than GREv0 and IPv4-in-IP. Sessions inside tunnels are not replicated and the new active node will not be able to reuse existing inspection verdicts to match the correct policy rules.
- Decrypted TLS/SSL connections—The decryption states are not synchronized, and if the active unit fails, then decrypted connections will be reset. New connections will need to be established to the new active unit. Connections that are not decrypted (in other words, those that match a TLS/SSL Do Not Decrypt rule action) are not affected and are replicated correctly.
- TCP state bypass connections
- Multicast routing.

Bridge Group Requirements for High Availability

There are special considerations for high availability when using bridge groups.

When the active unit fails over to the standby unit, the switch port running Spanning Tree Protocol (STP) can go into a blocking state for 30 to 50 seconds when it senses the topology change. To avoid traffic loss on the bridge group member interfaces while the port is in a blocking state, you can configure one of the following workarounds:

- Switch port is in Access mode—Enable the STP PortFast feature on the switch:

```
interface interface_id
  spanning-tree portfast
```

  The PortFast feature immediately transitions the port into STP forwarding mode upon linkup. The port still participates in STP. So if the port is to be a part of the loop, the port eventually transitions into STP blocking mode.

- If the switch port is in Trunk mode, or you cannot enable STP PortFast, then you can use one of the following less desirable workarounds that impacts failover functionality or STP stability:
  - Disable interface monitoring on the bridge group and member interfaces.
  - Increase the interface hold time in the failover criteria to a high value that will allow STP to converge before the unit fails over.
  - Decrease the STP timers on the switch to allow STP to converge faster than the interface hold time.

Failover Health Monitoring

The Firewall Threat Defense device monitors each unit for overall health and for interface health. This section includes information about how the Firewall Threat Defense device performs tests to determine the state of each unit.

Unit Health Monitoring

The Firewall Threat Defense device determines the health of the other unit by monitoring the failover link with hello messages. When a unit does not receive three consecutive hello messages on the failover link, the unit sends LANTEST messages on each data interface, including the failover link, to validate whether or not

the peer is responsive. The action that the Firewall Threat Defense device takes depends on the response from the other unit. See the following possible actions:

- If the Firewall Threat Defense device receives a response on the failover link, then it does not fail over.
- If the Firewall Threat Defense device does not receive a response on the failover link, but it does receive a response on a data interface, then the unit does not failover. The failover link is marked as failed. You should restore the failover link as soon as possible because the unit cannot fail over to the standby while the failover link is down.
- If the Firewall Threat Defense device does not receive a response on any interface, then the standby unit switches to active mode and classifies the other unit as failed.

Heartbeat Module Redundancy

Each unit in the HA periodically sends a broadcast keepalive heartbeat packet over the cluster control link. If the control plane is too busy handling traffic, sometimes the heartbeat packets do not reach the peers, or the peers do not process the heartbeat packets due to CPU overloading. When peers cannot communicate the keepalive status within the configurable timeout period, a false failover or split-brain scenario occurs.

The heartbeat module in the data plane helps to avoid the occurrence of false failover or split-brain due to traffic congestion in the control plane.

- The additional heartbeat module works similarly to the control plane module but sends and receives heartbeat messages using the data plane transport infrastructure.
- When the peer receives heartbeat packets in the data plane, a counter gets incremented.
- If the heartbeat transfer in the control plane fails, the node checks the heartbeat counter in the data plane. If the counter is incrementing, then the peer is alive, and the cluster does not perform a failover in this situation.

Note

- The additional heartbeat module is enabled by default whenever HA is enabled. You do not have to set a poll interval for the additional heartbeat module in the data plane. This module uses the same heartbeat interval that you set for the control plane.

Interface Monitoring

When a unit does not receive hello messages on a monitored interface for 15 seconds, it runs interface tests. If one of the interface tests fails for an interface, but this same interface on the other unit continues to successfully pass traffic, then the interface is considered to be failed, and the device stops running tests.

If the threshold you define for the number of failed interfaces is met (see **Devices** > **Device Management** > **High Availability** > **Failover Trigger Criteria**), and the active unit has more failed interfaces than the standby unit, then a failover occurs. If an interface fails on both units, then both interfaces go into the "Unknown" state and do not count towards the failover limit defined by failover interface policy.

An interface becomes operational again if it receives any traffic. A failed device returns to standby mode if the interface failure threshold is no longer met.

If an interface has IPv4 and IPv6 addresses configured on it, the device uses the IPv4 addresses to perform the health monitoring. If an interface has only IPv6 addresses configured on it, then the device uses IPv6

neighbor discovery instead of ARP to perform the health monitoring tests. For the broadcast ping test, the device uses the IPv6 all nodes address (FE02::1).

Interface Tests

The Firewall Threat Defense device uses the following interface tests. The duration of each test is approximately 1.5 seconds.

1. Link Up/Down test—A test of the interface status. If the Link Up/Down test indicates that the interface is down, then the device considers it failed, and testing stops. If the status is Up, then the device performs the Network Activity test.
2. Network Activity test—A received network activity test. At the start of the test, each unit clears its received packet count for its interfaces. As soon as a unit receives any eligible packets during the test, then the interface is considered operational. If both units receive traffic, then testing stops. If one unit receives traffic and the other unit does not, then the interface on the unit that does not receive traffic is considered failed, and testing stops. If neither unit receives traffic, then the device starts the ARP test.
3. ARP test—A test for successful ARP replies. Each unit sends a single ARP request for the IP address in the most recent entry in its ARP table. If the unit receives an ARP reply or other network traffic during the test, then the interface is considered operational. If the unit does not receive an ARP reply, then the device sends a single ARP request for the IP address in the *next* entry in the ARP table. If the unit receives an ARP reply or other network traffic during the test, then the interface is considered operational. If both units receive traffic, then testing stops. If one unit receives traffic, and the other unit does not, then the interface on the unit that does not receive traffic is considered failed, and testing stops. If neither unit receives traffic, then the device starts the Broadcast Ping test.
4. Broadcast Ping test—A test for successful ping replies. Each unit sends a broadcast ping, and then counts all received packets. If the unit receives any packets during the test, then the interface is considered operational. If both units receive traffic, then testing stops. If one unit receives traffic, and the other unit does not, then the interface on the unit that does not receive traffic is considered failed, and testing stops. If neither unit receives traffic, then testing starts over again with the ARP test. If both units continue to receive no traffic from the ARP and Broadcast Ping tests, then these tests will continue running in perpetuity.

Interface Status

Monitored interfaces can have the following status:

- Unknown—Initial status. This status can also mean the status cannot be determined.
- Normal—The interface is receiving traffic.
- Normal (Waiting)—The interface is up, but has not yet received a hello packet from the corresponding interface on the peer unit.
- Normal (Not-Monitored)—The interface is up, but is not monitored by the failover process.
- Testing—Hello messages are not heard on the interface for five poll times.
- Link Down—The interface or VLAN is administratively down.
- Link Down (Waiting)—The interface or VLAN is administratively down and has not yet received a hello packet from the corresponding interface on the peer unit.
- Link Down (Not-Monitored)—The interface or VLAN is administratively down, but is not monitored by the failover process.

- No Link—The physical link for the interface is down.
- No Link (Waiting)—The physical link for the interface is down and has not yet received a hello packet from the corresponding interface on the peer unit.
- No Link (Not-Monitored)—The physical link for the interface is down, but is not monitored by the failover process.
- Failed—No traffic is received on the interface, yet traffic is heard on the peer interface.

Failover Triggers and Detection Timing

The following events trigger failover in a Firepower high availability pair:

- More than 50% of the Snort instances on the active unit are down.
- Disk space on the active unit is more than 90% full.
- The **no failover active** command is run on the active unit or the **failover active** command is run on the standby unit.
- The active unit has more failed interfaces than the standby unit.
- Interface failure on the active device exceeds the threshold configured.

  By default, failure of a single interface causes failover. You can change the default value by configuring a threshold for the number of interfaces or a percentage of monitored interfaces that must fail for the failover to occur. If the threshold breaches on the active device, failover occurs. If the threshold breaches on the standby device, the unit moves to **Fail** state.

  To change the default failover criteria, enter the following command in global configuration mode:

  Table 28:

Command	Purpose
failover interface-policy num [%] `hostname (config)# failover interface-policy 20%`	Changes the default failover criteria. When specifying a specific number of interfaces, the *num* argument can be from 1 to 250. When specifying a percentage of interfaces, the *num* argument can be from 1 to 100.

The following table shows the failover triggering events and associated failure detection timing. If failover occurs, you can view the reason for the failover in the Message Center, along with various operations pertaining to the high availability pair. You can configure these thresholds to a value within the specified minimum-maximum range.

Table 29: Firewall Threat Defense Failover Times

Failover Triggering Event	Minimum	Default	Maximum
Active unit loses power, hardware goes down, or the software reloads or crashes. When any of these occur, the monitored interfaces or failover link do not receives any hello message.	800 milliseconds	15 seconds	45 seconds
Active unit interface physical link down.	500 milliseconds	5 seconds	15 seconds
Active unit interface up, but connection problem causes interface testing.	5 seconds	25 seconds	75 seconds

About Active/Standby Failover

Active/Standby failover lets you use a standby Firewall Threat Defense device to take over the functionality of a failed unit. When the active unit fails, the standby unit becomes the active unit.

Primary/Secondary Roles and Active/Standby Status

When setting up Active/Standby failover, you configure one unit to be primary and the other to be secondary. During configuration, the primary unit's policies are synchronized to the secondary unit. At this point, the two units act as a single device for device and policy configuration. However, for events, dashboards, reports and health monitoring, they continue to display as separate devices.

The main differences between the two units in a failover pair are related to which unit is active and which unit is standby, namely which IP addresses to use and which unit actively passes traffic.

However, a few differences exist between the units based on which unit is primary (as specified in the configuration) and which unit is secondary:

- The primary unit always becomes the active unit if both units start up at the same time (and are of equal operational health).
- The primary unit MAC addresses are always coupled with the active IP addresses. The exception to this rule occurs when the secondary unit becomes active and cannot obtain the primary unit MAC addresses over the failover link. In this case, the secondary unit MAC addresses are used.

Active Unit Determination at Startup

The active unit is determined by the following:

- If a unit boots and detects a peer already running as active, it becomes the standby unit.
- If a unit boots and does not detect a peer, it becomes the active unit.
- If both units boot simultaneously, then the primary unit becomes the active unit, and the secondary unit becomes the standby unit.

Failover Events

In Active/Standby failover, failover occurs on a unit basis.

The following table shows the failover action for each failure event. For each failure event, the table shows the failover policy (failover or no failover), the action taken by the active unit, the action taken by the standby unit, and any special notes about the failover condition and actions.

Table 30: Failover Events

Failure Event	Policy	Active Unit Action	Standby Unit Action	Notes
Active unit failed (power or hardware)	Failover	n/a	Become active Mark active as failed	No hello messages are received on any monitored interface or the failover link.
Formerly active unit recovers	No failover	Become standby	No action	None.
Standby unit failed (power or hardware)	No failover	Mark standby as failed	n/a	When the standby unit is marked as failed, then the active unit does not attempt to fail over, even if the interface failure threshold is surpassed.
Failover link failed during operation	No failover	Mark failover link as failed	Mark failover link as failed	You should restore the failover link as soon as possible because the unit cannot fail over to the standby unit while the failover link is down.
Failover link failed at startup	No failover	Become active Mark failover link as failed	Become active Mark failover link as failed	If the failover link is down at startup, both units become active.
State link failed	No failover	No action	No action	State information becomes out of date, and sessions are terminated if a failover occurs.
Interface failure on active unit above threshold	Failover	Mark active as failed	Become active	None.
Interface failure on standby unit above threshold	No failover	No action	Mark standby as failed	When the standby unit is marked as failed, then the active unit does not attempt to fail over even if the interface failure threshold is surpassed.

Config-Sync Optimization

When a device reboots or rejoins following a suspend or resume high availability, the joining device clears its running configuration. The active device then sends its entire configuration to the joining device for a full configuration synchronization. If the active device has a large configuration, this process can take several minutes.

The configuration sync optimization functionality enables comparing the configuration of the joining device and the active device by exchanging configuration hash values. If the hash computed on both active and

joining devices match, the joining device skips full configuration synchronization and rejoin the high availability configuration. This functionality ensures faster peering and reduces maintenance window and upgrade time.

Guidelines and Limitations of Config-Sync Optimization

- The configuration sync optimization functionality is enabled by default.
- Firewall Threat Defense multiple context mode supports configuration sync optimization by sharing the context order during full configuration synchronization, allowing comparison of context order during subsequent node-rejoin.
- If you configure passphrase and failover IPsec key, then configuration sync optimization is not effective as the hash value computed in the active and standby devices differs.
- If you configure the device with dynamic ACL or SNMPv3, configuration sync optimization is not effective.
- Active device synchronizes full configuration with flapping LAN links as default behavior. During failover flaps between active and standby devices, configuration sync optimization is not triggered and devices perform a full configuration synchronization.
- Configuration sync optimization gets triggered when the high availability configuration recovers from an interruption or loss of network communication between the active and standby devices.

Monitoring Config-Sync Optimization

When configuration sync optimization functionality is enabled, syslog messages are generated displaying whether the hash values computed on the active and joining unit match, does not match, or if the operation timeout expires. The syslog message also displays the time elapsed, from the time of sending the hash request to the time of getting and comparing the hash response.

Note The Configuration State in **show failover state** command displays the config sync state status during HA joining. This state does not reflect the later config deployments or changes, and replication on device until a resync has been initiated.

Requirements and Prerequisites for High Availability

Model Support

Secure Firewall Threat Defense

Supported Domains

Any

User Roles

Admin

Network Admin

Guidelines for High Availability

Model Support

- Firepower 1010 and Secure Firewall 1210/1220:
  - You should not use the switch port functionality when using High Availability. Because the switch ports operate in hardware, they continue to pass traffic on both the active *and* the standby units. High Availability is designed to prevent traffic from passing through the standby unit, but this feature does not extend to switch ports. In a normal High Availability network setup, active switch ports on both units will lead to network loops. We suggest that you use external switches for any switching capability. Note that VLAN interfaces can be monitored by failover, while switch ports cannot. Though Firewall Management Center allows you to configure switch ports / VLAN interfaces in HA, it may lead to misconfiguration due to network loops when connecting to external switches.
  - You can only use a firewall interface as the failover link.
- Firepower 9300—Intra-chassis High Availability is not supported.
- The Firewall Threat Defense Virtual on public cloud networks such as Microsoft Azure and Amazon Web Services are not supported with High Availability because Layer 2 connectivity is required.

Additional Guidelines

- When the active unit fails over to the standby unit, the connected switch port running Spanning Tree Protocol (STP) can go into a blocking state for 30 to 50 seconds when it senses the topology change. To avoid traffic loss while the port is in a blocking state, you can enable the STP PortFast feature on the switch:

  interface *interface_id* **spanning-tree portfast**

  This workaround applies to switches connected to both routed mode and bridge group interfaces. The PortFast feature immediately transitions the port into STP forwarding mode upon linkup. The port still participates in STP. So if the port is to be a part of the loop, the port eventually transitions into STP blocking mode.
- Configuring port security on the switches connected to the Firewall Threat Defense device failover pair can cause communication problems when a failover event occurs. This problem occurs when a secure MAC address configured or learned on one secure port moves to another secure port, a violation is flagged by the switch port security feature.
- For Active/Standby High Availability and a VPN IPsec tunnel, you cannot monitor both the active and standby units using SNMP over the VPN tunnel. The standby unit does not have an active VPN tunnel, and will drop traffic destined for the NMS. You can instead use SNMPv3 with encryption so the IPsec tunnel is not required.
- Both the peer devices go into unknown state and high-availability configuration fails if you run clish in any of the peer devices while creating a High Availability pair.
- Immediately after failover, the source address of syslog messages will be the failover interface address for a few seconds.

- For better convergence (during a failover), you must shut down the interfaces on a HA pair that are not associated with any configuration or instance.
- If you configure failover encryption in evaluation mode, the systems use DES for the encryption. If you then register the devices using an export-compliant account, the devices will use AES after a reboot. Thus, if a system reboots for any reason, including after installing an upgrade, the peers will be unable to communicate and both units will become the active unit. We recommend that you do not configure encryption until after you register the devices. If you do configure this in evaluation mode, we recommend you remove the encryption before registering the devices.
- When using SNMPv3 with failover, if you replace a failover unit, then SNMPv3 users are not replicated to the new unit. You must remove the users, re-add them, and then redeploy your configuration to force the users to replicate to the new unit.
- The device does not share SNMP client engine data with its peer.
- If you have a very large number of access control and NAT rules, the size of the configuration can prevent efficient configuration replication, resulting in the standby unit taking an excessively long time to reach standby ready state. This can also impact your ability to connect to the standby unit during replication through the console or SSH session. To enhance configuration replication performance, enable transactional commit for both access rules and NAT, using the **asp rule-engine transactional-commit access-group** and **asp rule-engine transactional-commit nat** commands.
- A unit in a High Availability pair transitioning to the standby role synchronizes its clock with the active unit.

  Example:

```
firepower#show clock
01:00:52 UTC Mar 1 2022

...
01:01:18 UTC Mar 1 2022 <======= Incorrect (previous) clock
Cold Standby              Sync Config               Detected an Active mate

19:38:21 UTC Apr 9 2022 <======= Updated clock
Sync Config               Sync File System          Detected an Active mate
...
firepower/sec/stby#show clock
19:38:40 UTC Apr 9 2022
```

- The units in High Availability do not dynamically synchronize the clock. Here are some examples of events when synchronization takes place:
  - A new High Availability pair is created.
  - High Availability is broken and re-created.
  - Communication over the failover link was disrupted and reestablished.
  - Failover status was manually changed at the CLI using the **no failover/failover** or **configure high-availability suspend/resume** (Firewall Threat Defense) commands.
- Enabling High Availability forces all routes to be deleted and are re-added after the High Availability progression changes to the Active state. You could experience connection loss during this phase.
- If you replace the primary unit, then when you re-create high-availability, you should set the replacement unit as the *secondary* unit so that the configurations are replicated from the former secondary unit to the

replacement unit. If you set the replacement unit as primary, you will accidentally overwrite the configuration that is present on the operational unit.

- Deploying Firepower 1100 devices in high availability with hundreds of interfaces configured on them can result in increased delay in the failover time (seconds).
- In the High Availability configuration, short-lived connections, usually using port 53, are closed quickly and never transferred or synchronized from Active to Standby, so there might be a difference in the number of connections on both High Availability devices. This is expected behavior for short-lived connections. You can try to compare the connections that are long-lived (for example, more than 30-60 seconds).
- In the High Availability configuration, embryonic connections—connection requests that have not yet completed the three-way handshake process—are closed quickly and not synchronized between the active and standby devices. This design ensures HA system efficiency and security. For this reason, there might be a difference in the number of connections on both High Availability devices, which is to be expected.
- If the failover LAN link is not connected back-to-back and instead connected through one or more switches, a failure within the intermediate path can cause the active unit to lose connectivity with the standby unit, resulting in inconsistent active/standby states. Although this does not impact High Availability functionality, it is recommended to check and recover the failover-link path between the active and standby units.

  When the failover LAN link is down, it is not recommended to deploy any configuration, as it may not be replicated to the peer unit.
- See the Cisco Secure Firewall Threat Defense Virtual Getting Started Guide and review your Firewall Threat Defense Virtual device configurations for high availability.
- In transparent mode, if you have problems with the active unit losing the MAC address of the hot-standby router (HSRP), create a static mapping for the MAC address.
- UCAPL or CC compliance mode cannot be changed if the threat defense device is in high availability. Modify the compliance mode before forming the high availability pair.

Add a High Availability Pair

When establishing an Active/Standby high-availability pair, you designate one of the devices as primary and the other as secondary. The Firewall Management Center deploys a merged configuration to the paired devices. If there is a conflict, the primary device setting is used.

In a multidomain deployment, devices in a high-availability pair must belong to the same domain.

Note The failover link and the stateful failover link are in a private IP space and are only used for communication between peers in a high-availability pair. After high availability is established, selected interface links and encryption settings cannot be modified without breaking the high-availability pair and reconfiguring it.

Caution Creating or breaking a high-availability pair immediately restarts the Snort process on the primary and secondary devices, temporarily interrupting traffic inspection on both devices. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information. The system warns you that continuing to create a high-availability pair restarts the Snort process on the primary and secondary devices and allows you to cancel.

Before you begin

Confirm that both devices:

- Are the same model.
- Have the same number and type of interfaces.
- Are in the same domain and group.
- Have normal health status and are running the same software.
- Are either in routed or transparent mode.

Note Only routed mode is supported for manager access on a data interface.

- Have the same NTP configuration. See Time Synchronization, on page 1002.
- Are fully deployed with no uncommitted changes.
- Do not have DHCP or PPPoE configured on any interfaces.
- For manager access on a data interface:
  - Use the same data interface on both devices for manager access.
  - You cannot use DHCP; only a static IP address is supported. Features that rely on DHCP cannot be used, including DDNS and zero-touch provisioning.

Note If you use zero-touch provisioning to register the device, when you use the outside interface for manager access, it uses DHCP by default. Before you can enable high availability, you need to change the IP address to a static address. See Change the Device IP Address, on page 158. Alternatively, you can use the Management interface instead; DHCP is supported on Management with high availability.

  - Have different static IP addresses in the same subnet.
  - Use the same manager configuration (**configure manager add** command) to ensure that the connectivity is the same.
  - You cannot use the data interface as the failover or state link.

Note The high availability formation is possible between the two Firewall Threat Defense devices when the certificate available on the primary device is not present on the secondary device. When high availability is formed, the certificate will be synched on the secondary device.

Procedure

Step 1 Add both devices to the Firewall Management Center according to Register With the Management Center, on page 31.

Step 2 Choose **Devices** > **Device Management**.

Step 3 From the **Add** drop-down menu, choose **High Availability**.

Step 4 Enter a display **Name** for the high-availability pair.

Step 5 Choose the **Primary Peer** device for the high-availability pair.

Step 6 Choose the **Secondary Peer** device for the high-availability pair.

Step 7 Click **Continue**.

Step 8 Under **LAN Failover Link**, choose an **Interface** with enough bandwidth to reserve for failover communications.

Note
Only interfaces that do not have a logical name and do not belong to a security zone, will be listed in the **Interface** drop-down in the **Add High Availability Pair** dialog.

Step 9 Type any identifying **Logical Name**.

Step 10 Type a **Primary IP** address for the failover link on the active unit.

This address should be on an unused subnet.

Note
169.254.1.0/24 and fd00:0:0:*::/64 are internally used subnets and cannot be used for the failover or state links.

Step 11 Optionally, choose **Use IPv6 Address**.

Step 12 Type a **Secondary IP** address for the failover link on the standby unit. This IP address must be in the same subnet as the primary IP address.

Step 13 If IPv4 addresses are used, type a **Subnet Mask** that applies to both the primary and secondary IP addresses.

Step 14 Optionally, under **Stateful Failover Link**, choose the same **Interface**, or choose a different interface and enter the high availability configuration information.

Note
169.254.1.0/24 and fd00:0:0:*::/64 are internally used subnets and cannot be used for the failover or state links.

Step 15 Optionally, choose **Enabled** and choose the **Key Generation** method for IPsec Encryption between the failover links.

Step 16 Click **OK**. This process takes a few minutes as the process synchronizes system data.

What to do next

Back up the devices. You can use the backup to quickly replace the devices when they fail and to restore the high availability service without being delinked from the Firewall Management Center. For more information, see Cisco Secure Firewall Management Center Administration Guide.

Configure Optional High Availability Parameters

You can view the initial High Availability Configuration on the Firewall Management Center. You cannot edit these settings without breaking the high availability pair and then re-establishing it.

You can edit the Failover Trigger Criteria to improve failover results. Interface Monitoring allows you to determine which interfaces are better suited for failover.

Configure Standby IP Addresses and Interface Monitoring

For each interface, set a standby IP address. Although recommended, the standby address is not required. Without a standby IP address, the active unit cannot perform network tests to check the standby interface health; it can only track the link state.

By default, monitoring is enabled on all physical interfaces, for the Firepower 1010 and Secure Firewall 1210/1220 all VLAN interfaces, with logical names configured. You might want to exclude interfaces attached to less critical networks from affecting your failover policy. Firepower 1010 and Secure Firewall 1210/1220 switch ports are not eligible for interface monitoring.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair you want to edit, click the **Edit** ( ).

Step 3 Click the **High Availability** tab.

Step 4 In the **Monitored Interfaces** area, click the **Edit** ( ) next to the interface you want to edit.

Step 5 Check the **Monitor this interface for failures** check box.

Step 6 On the **IPv4** tab, enter the **Standby IP Address**.

This address must be a free address on the same network as the active IP address.

Step 7 If you configured the IPv6 address manually, on the **IPv6** tab, click the **Edit** ( ) next to the active IP address, enter the **Standby IP Address**, and click **OK**.

This address must be a free address on the same network as the active IP address. For autogenerated and **Enforce EUI 64** addresses, the standby address is automatically generated.

Step 8 Click **OK**.

Edit High Availability Failover Criteria

You can customize failover criteria based on your network deployment.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair you want to edit, click the **Edit** ( ).

Step 3 Choose **High Availability**.

Step 4 Next to **Failover Trigger Criteria**, click the **Edit** ( ).

Step 5 Under **Interface Failure Threshold**, choose the number or percentage of interfaces that must fail before the device fails over.

Step 6 Under **Hello packet Intervals**, choose how often hello packets are sent over the failover link.

Step 7 Click **OK**.

Configure Virtual MAC Addresses

You can configure active and standby MAC addresses for failover using the following methods in the Secure Firewall Management Center:

- From the **Advanced tab** on the **Edit Interface** page during interface configuration; see Configure the MAC Address, on page 881.
- From the **Add Interface MAC Address** dialog-box which is accessed from the **High Availability** page; see this procedure.

Note To configure the MAC address in both primary and secondary units (so that the MAC address is transferred to all sub-interfaces to both the high-availability units), the recommended approach is to use the **Interfaces** tab to replicate the MAC addresses on sub-interfaces over both active and standby high-availability units.

If you configure active and standby MAC addresses in both locations, the addresses defined during interface configuration take precedence for failover.

You can minimize loss of traffic during failover by designating active and standby MAC addresses to the physical interface. This feature offers redundancy against IP address mapping for failover.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair you want to edit, click **Edit** ( ).

Step 3 Click **High Availability**.

Step 4 Click the **Add** ( ) icon next to Interface MAC Addresses.

Step 5 Choose a **Physical Interface**.

Step 6 Enter the **Active Interface Mac Address**.

Step 7 Enter the **Standby Interface Mac Address**.

Step 8 Click **OK**.

Note
For detailed information, see Task 2, steps from 10 to 14 in Configure FTD High Availability on Firepower Appliances.

.

Manage High Availability

This section describes how to manage High Availability units after you enable High Availability, including how to change the High Availability setup and how to force failover from one unit to another.

Switch the Active Peer in the Firewall Threat Defense High Availability Pair

After you establish the Firewall Threat Defense high availability pair, you can manually switch the active and standby units, effectively forcing failover for reasons such as persistent fault or health events on the current active unit. Both units should be fully deployed before you complete this procedure.

Before you begin

Refresh Node Status for a Single Firewall Threat Defense High Availability Pair, on page 446. This ensures that the status on the Firewall Threat Defense high availability device pair is in sync with the status on the Firewall Management Center.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the high availability pair where you want to change the active peer, click the **Switch Active Peer**.

Step 3 You can:

- Click **Yes** to immediately make the standby device the active device in the high availability pair.
- Click **No** to cancel and return to the Device Management page.

Refresh Node Status for a Single Firewall Threat Defense High Availability Pair

Whenever active or standby devices in the Firewall Threat Defense high availability pair are rebooted, the Firewall Management Center may not display accurate high availability status for either device. This is because when the device reboots, the high availability status is immediately updated on the device and its corresponding event is sent to the Firewall Management Center. However, the status may not be updated on the Firewall

Management Center because the communication between the device and the Firewall Management Center is yet to be established.

Communication failures or weak communication channels between the Firewall Management Center and devices may result in out of sync data. When you switch the active and standby devices in a high availability pair, the change may not be reflected in the Firewall Management Center even after a significant time duration.

In these scenarios, you can refresh the high availability node status to obtain accurate information about the active and standby device in a high availability pair.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the high availability pair where you want to refresh the node status, click the **Refresh HA Node Status**.

Step 3 Click **Yes** to refresh the node status.

Suspend and Resume High Availability

You can suspend a unit in a high-availability pair, which is useful when:

- Both units are in an active-active situation, and fixing the communication on the failover link does not correct the problem.
- You want to troubleshoot an active or standby unit and do not want the units to fail over during that time.

When you suspend high availability, the currently active device remains active, handling all user connections. However, failover criteria are no longer monitored, and the system will never fail over to the now pseudo-standby device.

When using a data interface for manager access, the management connection will be broken until after you resume.

The key difference between suspending high availability and breaking high availability is that on a suspended high-availability device, the high-availability configuration is retained. When you break high availability, the configuration is erased. Thus, you have the option to resume high availability on a suspended system, which enables the existing configuration and makes the two devices function as a failover pair again.

To suspend high availability, use the **configure high-availability suspend** command.

```
> configure high-availability suspend
Please ensure that no deployment operation is in progress before suspending
high-availability.
Please enter 'YES' to continue if there is no deployment operation in
progress and 'NO' if you wish to abort: YES
Successfully suspended high-availability.
```

If you suspend high availability from the active unit, the configuration is suspended on both the active and standby units. The standby unit interface configuration is also erased. If you suspend it from the standby unit, it is suspended on the standby unit only, but the active unit will not attempt to fail over to a suspended unit.

To resume failover, use the **configure high-availability resume** command.

```
> configure high-availability resume
Successfully resumed high-availablity.
```

You can resume a unit only if it is in Suspended state. The unit will negotiate active/standby status with the peer unit.

Note Suspending high availability is a temporary state. If you reload a unit, it resumes the high-availability configuration automatically and negotiates the active/standby state with the peer.

Replace a Unit in Firewall Threat Defense High Availability Pair

To replace a failed unit in the Firewall Threat Defense high availability pair using a backup file, see *Restoring Firewall Management Centers and Managed Devices* in the Cisco Secure Firewall Management Center Administration Guide.

If you do not have a backup of the failed device, you must break high availability. Then, register the replacement device to the Secure Firewall Management Center and reestablish high availability. The process varies depending on whether the device is primary or secondary:

- Replace a Primary Firewall Threat Defense HA Unit with no Backup, on page 448
- Replace a Secondary Firewall Threat Defense HA Unit with no Backup, on page 449

Replace a Primary Firewall Threat Defense HA Unit with no Backup

Follow the steps below to replace a failed primary unit in the Firewall Threat Defense high availability pair. Failing to follow these steps can overwrite the existing high availability configuration.

Caution Creating or breaking the Firewall Threat Defense high availability pair immediately restarts the Snort process on the primary and secondary devices, temporarily interrupting traffic inspection on both devices. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information. The system warns you that continuing to create a high availability pair restarts the Snort process on the primary and secondary devices and allows you to cancel.

Caution Never move a disk from sensor or Firewall Management Center to another device without reimaging the disk. This is an unsupported configuration and can cause breakage in functionality.

Procedure

Step 1 Choose **Force Break** to separate the high availability pair; see Break a High Availability Pair, on page 450.

Note
The break operation removes all the configuration related to HA from Firewall Threat Defense and Firewall Management Center, and you need to recreate it manually later. To successfully configure the same HA pair, ensure that you save the IPs, MAC addresses, and monitoring configuration of all the interfaces/subinterfaces prior to executing the HA break operation.

Step 2 Unregister the failed primary Firewall Threat Defense device from the Firewall Management Center; see Unregister a Device from the Firewall Management Center, on page 60.

Step 3 Register the replacement Firewall Threat Defense to the Firewall Management Center; see Register With the Management Center, on page 31.

Step 4 Configure high availability, using the existing secondary/active unit as the primary device and the replacement device as the secondary/standby device during registration; see Add a High Availability Pair, on page 441.

Replace a Secondary Firewall Threat Defense HA Unit with no Backup

Follow the steps below to replace a failed secondary unit in the Firewall Threat Defense high availability pair.

Caution Creating or breaking the Firewall Threat Defense high availability pair immediately restarts the Snort process on the primary and secondary devices, temporarily interrupting traffic inspection on both devices. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information. The system warns you that continuing to create a high availability pair restarts the Snort process on the primary and secondary devices and allows you to cancel.

Procedure

Step 1 Choose **Force Break** to separate the high availability pair; see Break a High Availability Pair, on page 450.

Note
The break operation removes all the configuration related to HA from Firewall Threat Defense and Firewall Management Center, and you need to recreate it manually later. To successfully configure the same HA pair, ensure that you save the IPs, MAC addresses, and monitoring configuration of all the interfaces/subinterfaces prior to executing the HA break operation.

Step 2 Unregister the secondary Firewall Threat Defense device from the Firewall Management Center; see Unregister a Device from the Firewall Management Center, on page 60.

Step 3 Register the replacement Firewall Threat Defense to the Firewall Management Center; see Register With the Management Center, on page 31.

Step 4 Configure high availability, using the existing primary/active unit as the primary device and the replacement device as the secondary/standby device during registration; see Add a High Availability Pair, on page 441.

Break a High Availability Pair

When you break a high-availability pair, the high-availability configuration is removed from both units.

When using the Management interface for manager access: The active unit remains up and passing traffic. The standby unit interface configuration is erased.

When using a data interface for manager access: See the following details.

- The active unit remains up and passing traffic.
- The standby unit data interfaces are shut down except for the manager access interface, which remains up using the standby IP address so it can maintain the management connection.
- In a remote branch deployment setup, all standby unit data interfaces that are assigned with a logical name are shut down except for the manager access interface, which remains up to maintain the management connection.
- If the primary unit is in the standby state:
  - The IP addresses for manager access are swapped permanently in the Firewall Management Center configuration: the primary unit uses the standby IP address, and the secondary unit uses the active IP address.
  - If the Firewall Management Center initiated the management connection, and you specified a hostname for the device, then you need to update the DNS server so the swapped IP addresses are associated with the correct hostnames.
  - Breaking high availability causes a deployment to the standby unit. If the management connection is not yet reestablished because of the swapped IP addresses, the deployment may fail. In this case, you will need to manually trigger the deployment later (after the management connection is established). Be sure to complete the standby deployment before deploying changes to the active unit.

Policies that were not deployed to the active unit prior to the break operation continue to remain un-deployed after the break operation is complete. Deploy the policies on the standalone device after the break operation is complete.

Note

- When IPsec is enabled on high availability interfaces on the Firewall Threat Defense device, the device cannot prioritize the encrypted packets into the high-priority receive queue. As a result, during high-volume data traffic scenarios, attempts to break high availability may fail as the device cannot efficiently manage and prioritize the large number of encrypted connections. To view the device's resource usage and the maximum throughput, use the `show resource usage` command.
- If you cannot reach the high-availability pair using the Firewall Management Center, connect to the CLI on each device and enter **configure high-availability disable** to manually break high availability. See also Unregister a High Availability Pair and Register to a New Firewall Management Center, on page 451.

Caution Breaking the Firewall Threat Defense high-availability pair immediately restarts the Snort process on the primary and secondary units, temporarily interrupting traffic inspection on both devices. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

Before you begin

- Refresh Node Status for a Single Firewall Threat Defense High Availability Pair, on page 446. This ensures that the status on the high-availability pair is in sync with the status on the Firewall Management Center.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the high-availability pair you want to break, click **More** (⋮) and choose **Break**.

Step 3 If the standby peer does not respond, check **Force Break**.

Step 4 Click **Yes**.

The Break operation removes the high-availability configuration from the active and standby units.

A FlexConfig policy deployed on the active unit may show a deployment failure after the break high-availability operation. You must alter and re-deploy the FlexConfig policy on the active unit.

What to do next

If you are using a FlexConfig policy on the active unit, alter and re-deploy the FlexConfig policy to eliminate deployment errors.

Note After you break high availability, the Firewall Threat Defense device that was operating as the active unit will still have the standby unit's IP address listed in its configuration. To resolve this, do an additional deployment on the formerly active Firewall Threat Defense device to remove the standby unit's IP address from its configuration.

Unregister a High Availability Pair and Register to a New Firewall Management Center

You can unregister the pair from the Firewall Management Center, which keeps the High Availability pair intact. You might want to unregister the pair if you want to register it to a new Firewall Management Center or if the Firewall Management Center can no longer reach the pair.

Unregistering a High Availability pair:

- Severs all communication between the Firewall Management Center and the pair.
- Removes the pair from the **Device Management** page.
- Returns the pair to local time management if the pair's platform settings policy is configured to receive time from the Firewall Management Center using NTP.
- Leaves the configuration intact, so the pair continues to process traffic.

  Policies, such as NAT and VPN, ACLs, and the interface configurations remain intact.

Registering the pair again to the same or a different Firewall Management Center causes the configuration to be removed, so the pair will stop processing traffic at that point; the High Availability configuration remains intact so you can add the pair as a whole. You can choose an access control policy at registration, but you will have to re-apply other policies after registration and then deploy the configuration before it will process traffic again.

Before you begin

- This procedure requires CLI access to the primary unit.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the High Availability pair you want to unregister, click **More** (⋮) and choose **Unregister**.

Step 3 Click **Yes**. The device High Availability pair is unregistered.

Step 4 You can register the pair to a new (or the same) Firewall Management Center by adding the primary unit as a new device.

a) Determine the primary unit by connecting to the CLI on one unit, and entering the **show failover** command.

The first line of the output shows whether this unit is Primary or Secondary.

```
> show failover
Failover unit Primary
Failover LAN Interface: failover GigabitEthernet0/2 (up)
Reconnect timeout 0:00:00
Unit Poll frequency 1 seconds, holdtime 15 seconds
Failover On

[...]
```

b) At the primary unit's CLI, identify the new Firewall Management Center using the **configure manager add** command. See Modify Firewall Threat Defense Management Interfaces at the CLI, on page 158.

c) Choose **Devices** > **Device Management**, and then click **Add** > **Device**.

You only need to add the primary unit as a device, and the Firewall Management Center will discover the secondary unit.

Monitoring High Availability

This section lets you monitor the High Availability status.

View Failover History

You can view the failover history of both high availability devices in a single view. The history displays in chronological order and includes the reason for any failover.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair you want to edit, click **Edit** ( ).

Step 3 Choose **Summary**.

Step 4 Under General, click **View** ( ).

View Stateful Failover Statistics

You can view the stateful failover link statistics of both the primary and secondary devices in the high availability pair.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the device high-availability pair you want to edit, click **Edit** ( ).

Step 3 Choose **High Availability**.

Step 4 Under Stateful Failover Link, click **View** ( ).

Step 5 Choose a device to view statistics.

Troubleshoot Configuration Sync Failure

When forming a failover pair, the joining unit clears its running configuration and replicates the entire configuration from the active unit. Upon completing full configuration sync, the joining unit assumes the standby ready role and establishes the failover pair. After the unit joins the failover pair, any configuration change on the active unit are also replicated on the standby unit to keep both the units synchronized.

If the standby unit fails to replicate any configuration change commands, it reports configuration sync failure and exit the high availability by disabling failover. This section describes steps for identifying and troubleshooting configuration sync failure error reported by the standby unit.

To view the configuration sync errors or stats, you can use the following CLI commands through an SSH session or the Threat Defense CLI:

- **show failover config-sync errors all** to display all configuration synchronization errors related to failover.
- **show failover config-sync stats all** to view statistics regarding failover configuration synchronization.

To re-enable high availability:

- Re-enable failover by executing the **failover reset** command on the active unit.
- If re-enabling the failover is not successful, delete or update the configuration change that the standby unit failed to replicate, and then re-enable the failover.

History for High Availability

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Improved role switch time during failover	7.7	7.7	When a failover occurs, the new active device generates multicast packets for each MAC address entry and sends them to all bridge group interfaces, prompting the upstream switches to update their routing tables. This task of generating and sending multicast packets to the bridge interfaces now runs asynchronously in the data plane, allowing critical failover tasks in the control plane to proceed without delays. This enhancement improves role switch time during a failover and reduces downtime.
High availability support for the manager access data interface	7.4	7.4	You can now use a data interface for manager access with Firewall Threat Defense high availability.
Unregistering a high-availability pair now allows you to re-register without breaking the pair	7.3	Any	When you delete (unregister) a high-availability pair, you no longer have to manually break the pair at the CLI and re-register standalone devices. You can now add the primary unit to a new Firewall Management Center, and the standby unit will be discovered automatically. Re-registering the pair will still erase the configuration, and your policies will need to be re-applied.
Policy rollback support for high availability	7.2	Any	The **configure policy rollback** command is supported for high availability.
Config-Sync Optimization feature for faster HA peering	7.2	Any	The Config-Sync Optimization feature enables comparing the configuration of the joining unit and the active unit by exchanging config-hash values. If the hash computed on both active and joining units match, the joining unit skips full config-sync and rejoin the HA. This feature enables faster HA peering and reduces maintenance window and upgrade time.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Improvements to the upgrade workflow for clustered and high-availability devices	7.1	Any	We made the following improvements to the upgrade workflow for clustered and high-availability devices: • The upgrade wizard now correctly displays clustered and high-availability units as groups, rather than as individual devices. The system can identify, report, and preemptively require fixes for group-related issues you might have. For example, you cannot upgrade a cluster on the Firepower 4100/9300 if you have made unsynced changes on Firepower Chassis Manager. • We improved the speed and efficiency of copying upgrade packages to clusters and high-availability pairs. Previously, the FMC copied the package to each group member sequentially. Now, group members can get the package from each other as part of their normal sync process. • You can now specify the upgrade order of data units in a cluster. The control unit always upgrades last.
Clearing routes in a high-availability group or cluster.	7.1	Any	In previous releases, the **clear route** command cleared the routing table on the unit only. Now, when operating in a high-availability group or cluster, the command is available on the active or control unit only, and clears the routing table on all units in the group or cluster.

FTD High Availability Hardening	6.2.3	Any	Version 6.2.3 introduces the following features for FTD devices in high availability:
			• Whenever active or standby FTD devices in a high-availability pair restart, the FMC may not display accurate high-availability status for either managed device. However, the status may not upgrade on the FMC because the communication between the device and the FMC is not established yet. The **Refresh Node Status** option on the **Devices** > **Device Management** page allows you to refresh the high-availability unit status to obtain accurate information about the active and standby device in a high-availability pair.
			• The **Devices** > **Device Management** page of the FMC UI has a new **Switch Active Peer** icon.
			• Version 6.2.3 includes a new REST API object, **Device High Availability Pair Services**, that contains four functions:
			  • **DELETE ftddevicehapairs**
			  • **PUT ftddevicehapairs**
			  • **POST ftddevicehapairs**
			  • **GET ftddevicehapairs**

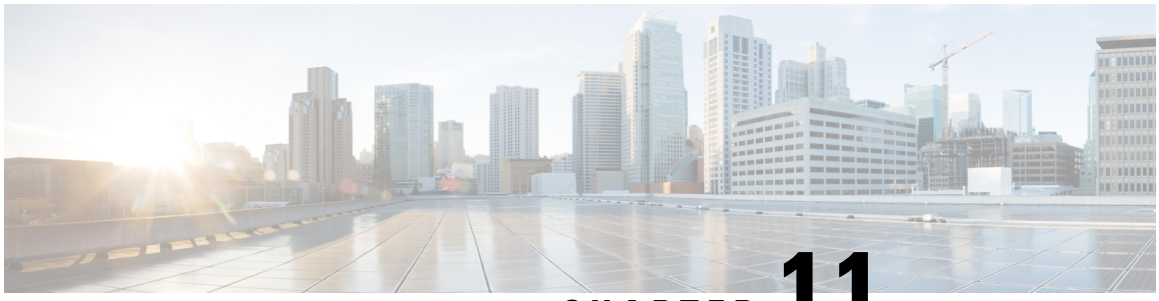

CHAPTER 11

Clustering for the Secure Firewall 3100/4200

Clustering lets you group multiple Firewall Threat Defense units together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices.

Note Some features are not supported when using clustering. See Unsupported Features with Clustering, on page 509.

- About Clustering for the Secure Firewall 3100/4200, on page 457
- Licenses for Clustering, on page 459
- Requirements and Prerequisites for Clustering, on page 459
- Guidelines for Clustering, on page 460
- Configure Clustering, on page 464
- Manage Cluster Nodes, on page 489
- Monitoring the Cluster, on page 499
- Troubleshooting the Cluster, on page 505
- Examples for Clustering, on page 507
- Reference for Clustering, on page 509
- History for Clustering, on page 522

About Clustering for the Secure Firewall 3100/4200

This section describes the clustering architecture and how it works.

How the Cluster Fits into Your Network

The cluster consists of multiple firewalls acting as a single unit. To act as a cluster, the firewalls need the following infrastructure:

- Isolated, high-speed backplane network for intra-cluster communication, known as the *cluster control link*.
- Management access to each firewall for configuration and monitoring.

When you place the cluster in your network, the upstream and downstream routers need to be able to load-balance the data coming to and from the cluster using one of the following methods:

- Spanned EtherChannel (Recommended)—Interfaces on multiple members of the cluster are grouped into a single EtherChannel; the EtherChannel performs load balancing between units.
- Policy-Based Routing (Routed firewall mode only)—The upstream and downstream routers perform load balancing between units using route maps and ACLs.
- Equal-Cost Multi-Path Routing (Routed firewall mode only)—The upstream and downstream routers perform load balancing between units using equal cost static or dynamic routes.

Control and Data Node Roles

One member of the cluster is the control node. If multiple cluster nodes come online at the same time, the control node is determined by the priority setting; the priority is set between 1 and 100, where 1 is the highest priority. All other members are data nodes. When you first create the cluster, you specify which node you want to be the control node, and it will become the control node simply because it is the first node added to the cluster.

All nodes in the cluster share the same configuration. The node that you initially specify as the control node will overwrite the configuration on the data nodes when they join the cluster, so you only need to perform initial configuration on the control node before you form the cluster.

Some features do not scale in a cluster, and the control node handles all traffic for those features.

Cluster Interfaces

You can configure data interfaces as either Spanned EtherChannels or as Individual interfaces. All data interfaces in the cluster must be one type only. See About Cluster Interfaces, on page 464 for more information.

For Spanned EtherChannels: You can use regular firewall interfaces or IPS-only interfaces (inline sets or passive interfaces). For Individual interfaces: IPS-only interfaces are not supported.

Cluster Control Link

Each unit must dedicate at least one hardware interface as the cluster control link. See Cluster Control Link, on page 464 for more information.

Configuration Replication

All nodes in the cluster share a single configuration. You can only make configuration changes on the control node (with the exception of the bootstrap configuration), and changes are automatically synced to all other nodes in the cluster.

Management Network

You must manage each node using the Management interface; management from a data interface is not supported with clustering.

Licenses for Clustering

You assign feature licenses to the cluster as a whole, not to individual nodes. However, each node of the cluster consumes a separate license for each feature. The clustering feature itself does not require any licenses.

When you add the control node to the Firewall Management Center, you can specify the feature licenses you want to use for the cluster. Before you create the cluster, it doesn't matter which licenses are assigned to the data nodes; the license settings for the control node are replicated to each of the data nodes. You can modify licenses for the cluster by clicking **Edit Licenses** in **System** (⚙) > **Licenses** > **Smart Licenses** or choosing**Devices** > **Device Management**, clicking **Edit** (✎) for the cluster, and then in the **Cluster** > **License** area, clicking **Edit** (✎).

Note If you add the cluster before the Firewall Management Center is licensed (and running in Evaluation mode), then when you license the Firewall Management Center, you can experience traffic disruption when you deploy policy changes to the cluster. Changing to licensed mode causes all data units to leave the cluster and then rejoin.

Requirements and Prerequisites for Clustering

Model Requirements

- Secure Firewall 3100—Maximum 16 units
- Secure Firewall 4200—Maximum 16 units

User Roles

- Admin
- Access Admin
- Network Admin

Hardware and Software Requirements

All units in a cluster:

- Must be the same model.
- Must include the same interfaces.
- The Firewall Management Center access must be from the Management interface; data interface management is not supported.
- Must run the identical software except at the time of an image upgrade. Hitless upgrade is supported.
- Must be in the same firewall mode, routed or transparent.
- Must be in the same domain.

- Must be in the same group.
- Must not have any deployment pending or in progress.
- The control node must not have any unsupported features configured (see Unsupported Features with Clustering, on page 509).
- Data nodes must not have any VPN configured. The control node can have site-to-site VPN configured.

Switch Requirements

- Be sure to complete the switch configuration before you configure clustering. Make sure the ports connected to the cluster control link have the correct (higher) MTU configured. By default, the cluster control link MTU is set to 100 bytes higher than the data interfaces. If the switches have an MTU mismatch, the cluster formation will fail.

Guidelines for Clustering

Firewall Mode

The firewall mode must match on all units.

High Availability

High Availability is not supported with clustering.

IPv6

The cluster control link is only supported using IPv4.

Switches

- Make sure connected switches match the MTU for both cluster data interfaces and the cluster control link interface. You should configure the cluster control link interface MTU to be at least 100 bytes higher than the data interface MTU, so make sure to configure the cluster control link connecting switch appropriately. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead. In addition, we do not recommend setting the cluster control link MTU between 2561 and 8362; due to block pool handling, this MTU size is not optimal for system operation. When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the ping fails, a notification is generated so you can fix the MTU mismatch on connecting switches and try again.
- For Cisco IOS XR systems, if you want to set a non-default MTU, set the IOS XR interface MTU to be 14 bytes higher than the cluster device MTU. Otherwise, OSPF adjacency peering attempts may fail unless the **mtu-ignore** option is used. Note that the cluster device MTU should match the IOS XR *IPv4* MTU. This adjustment is not required for Cisco Catalyst and Cisco Nexus switches.
- On the switch(es) for the cluster control link interfaces, you can optionally enable Spanning Tree PortFast on the switch ports connected to the cluster unit to speed up the join process for new units.

- On the switch, we recommend that you use one of the following EtherChannel load-balancing algorithms: **source-dest-ip** or **src-dst-mixed-ip-port** (see the Cisco Nexus OS and Cisco IOS-XE **port-channel load-balance** command). Do not use a **vlan** keyword in the load-balance algorithm because it can cause unevenly distributed traffic to the devices in a cluster.
- If you change the load-balancing algorithm of the EtherChannel on the switch, the EtherChannel interface on the switch temporarily stops forwarding traffic, and the Spanning Tree Protocol restarts. There will be a delay before traffic starts flowing again.
- Switches on the cluster control link path should not verify the L4 checksum. Redirected traffic over the cluster control link does not have a correct L4 checksum. Switches that verify the L4 checksum could cause traffic to be dropped.
- Port-channel bundling downtime should not exceed the configured keepalive interval.
- On Supervisor 2T EtherChannels, the default hash distribution algorithm is adaptive. To avoid asymmetric traffic in a VSS design, change the hash algorithm on the port-channel connected to the cluster device to fixed:

  router(config)# **port-channel** *id* **hash-distribution fixed**

  Do not change the algorithm globally; you may want to take advantage of the adaptive algorithm for the VSS peer link.

- You should disable the LACP Graceful Convergence feature on all cluster-facing EtherChannel interfaces for Cisco Nexus switches.

EtherChannels

- In Catalyst 3750-X Cisco IOS software versions earlier than 15.1(1)S2, the cluster unit did not support connecting an EtherChannel to a switch stack. With default switch settings, if the cluster unit EtherChannel is connected cross stack, and if the control unit switch is powered down, then the EtherChannel connected to the remaining switch will not come up. To improve compatibility, set the **stack-mac persistent timer** command to a large enough value to account for reload time; for example, 8 minutes or 0 for indefinite. Or, you can upgrade to more a more stable switch software version, such as 15.1(1)S2.
- Spanned vs. Device-Local EtherChannel Configuration—Be sure to configure the switch appropriately for Spanned EtherChannels vs. Device-local EtherChannels.
  - Spanned EtherChannels—For cluster unit *Spanned* EtherChannels, which span across all members of the cluster, the interfaces are combined into a single EtherChannel on the switch. Make sure each interface is in the same channel group on the switch.

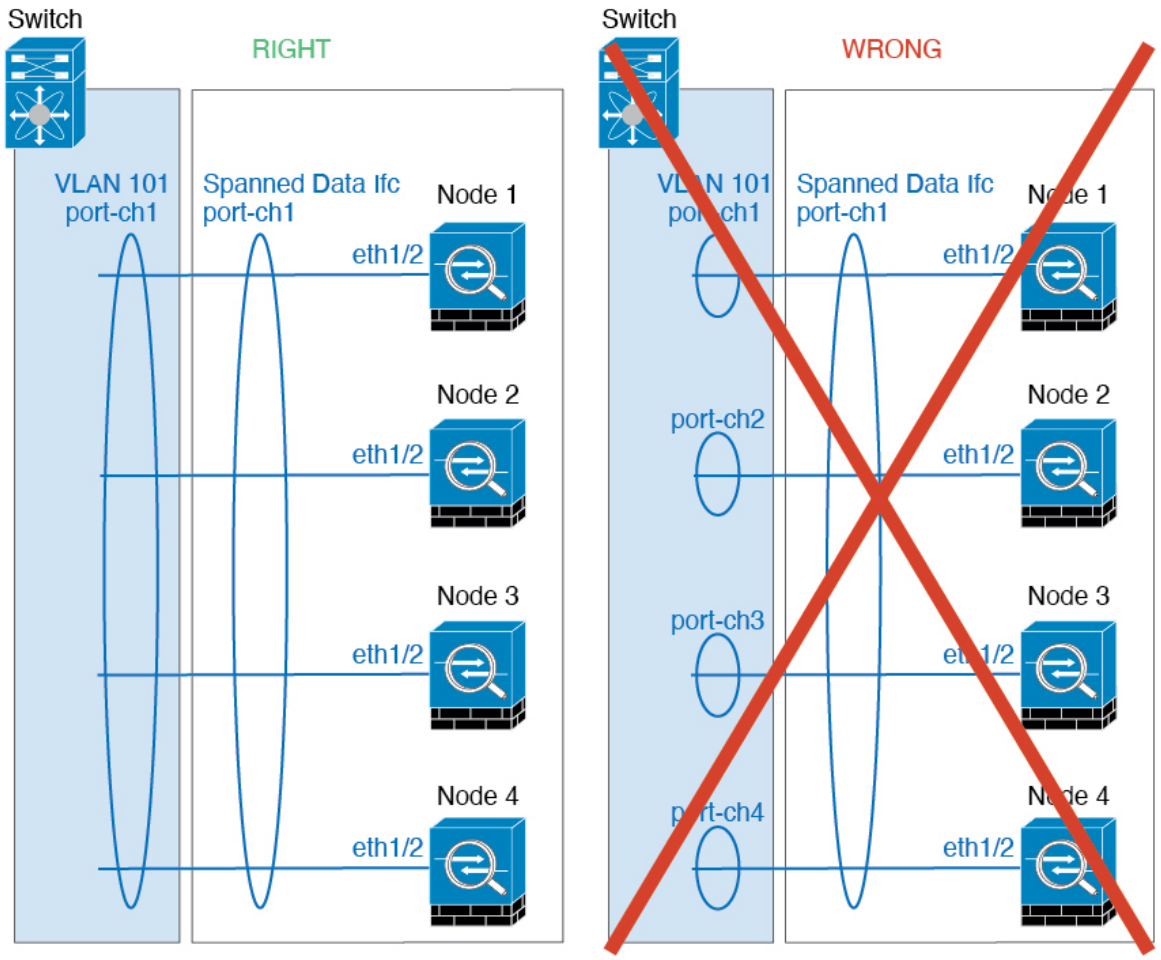

- Device-local EtherChannels—For cluster unit *Device-local* EtherChannels including any EtherChannels configured for the cluster control link, be sure to configure discrete EtherChannels on the switch; do not combine multiple cluster unit EtherChannels into one EtherChannel on the switch.

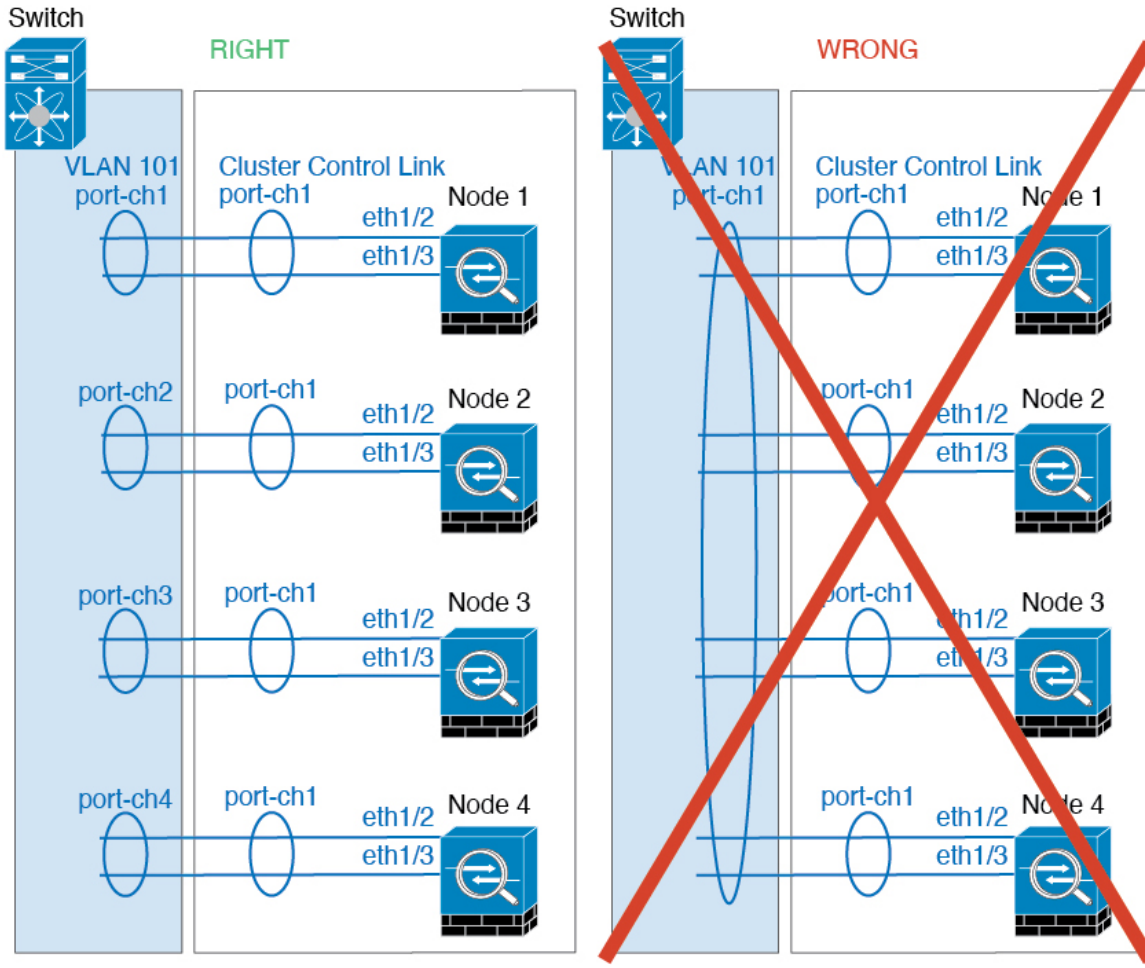

Additional Guidelines

- When significant topology changes occur (such as adding or removing an EtherChannel interface, enabling or disabling an interface on the Firewall Threat Defense or the switch, adding an additional switch to form a VSS or vPC) you should disable the health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all units, you can re-enable the interface health check feature.
- When adding a unit to an existing cluster, or when reloading a unit, there will be a temporary, limited packet/connection drop; this is expected behavior. In some cases, the dropped packets can hang your connection; for example, dropping a FIN/ACK packet for an FTP connection will make the FTP client hang. In this case, you need to reestablish the FTP connection.
- If you use a Windows 2003 server connected to a Spanned EtherChannel, when the syslog server port is down and the server does not throttle ICMP error messages, then large numbers of ICMP messages are sent back to the ASA cluster. These messages can result in some units of the ASA cluster experiencing high CPU, which can affect performance. We recommend that you throttle ICMP error messages.
- For decrypted TLS/SSL connections, the decryption states are not synchronized, and if the connection owner fails, then decrypted connections will be reset. New connections will need to be established to a

new unit. Connections that are not decrypted (they match a do-not-decrypt rule) are not affected and are replicated correctly.

Defaults for Clustering

- The cLACP system ID is auto-generated, and the system priority is 1 by default.
- The cluster health check feature is enabled by default with the holdtime of 3 seconds. Interface health monitoring is enabled on all interfaces by default.
- The cluster auto-rejoin feature for a failed cluster control link is unlimited attempts every 5 minutes.
- The cluster auto-rejoin feature for a failed data interface is 3 attempts every 5 minutes, with the increasing interval set to 2.
- Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

Configure Clustering

To add a cluster to the Firewall Management Center, add each node to the Firewall Management Center as a standalone unit, configure interfaces on the unit you want to make the control node, and then form the cluster.

About Cluster Interfaces

You can configure data interfaces as either Spanned EtherChannels or as Individual interfaces. All data interfaces in the cluster must be one type only. You cannot configure Ethernet 1/1 as a Spanned EtherChannel and configure Ethernet 1/2 as an Individual interface within the same cluster, for example.

For Spanned EtherChannels: You can use regular firewall interfaces or IPS-only interfaces (inline sets or passive interfaces). For Individual interfaces: IPS-only interfaces are not supported.

Each unit must also dedicate at least one hardware interface as the cluster control link.

Cluster Control Link

Each unit must dedicate at least one hardware interface as the cluster control link. We recommend using an EtherChannel for the cluster control link if available.

Cluster Control Link Traffic Overview

Cluster control link traffic includes both control and data traffic.

Control traffic includes:

- Control node election.
- Configuration replication.
- Health monitoring.

Data traffic includes:

- State replication.

- Connection ownership queries and data packet forwarding.

Cluster Control Link Interfaces and Network

You can use any physical interface or EtherChannel for the cluster control link. You cannot use a VLAN subinterface as the cluster control link. You also cannot use the Management interface.

Each cluster control link has an IP address on the same subnet. This subnet should be isolated from all other traffic, and should include only the cluster control link interfaces.

Note For a 2-member cluster, do not directly-connect the cluster control link from one node to the other node. If you directly connect the interfaces, then when one unit fails, the cluster control link fails, and thus the remaining healthy unit fails. If you connect the cluster control link through a switch, then the cluster control link remains up for the healthy unit. If you need to directly-connect the units (for testing purposes, for example), then you should configure and enable the cluster control link interface on both nodes before you form the cluster.

Size the Cluster Control Link

If possible, you should size the cluster control link to match the expected throughput of each chassis so the cluster control link can handle the worst-case scenarios.

Cluster control link traffic is comprised mainly of state update and forwarded packets. The amount of traffic at any given time on the cluster control link varies. The amount of forwarded traffic depends on the load-balancing efficacy or whether there is a lot of traffic for centralized features. For example:

- NAT results in poor load balancing of connections, and the need to rebalance all returning traffic to the correct units.
- When membership changes, the cluster needs to rebalance a large number of connections, thus temporarily using a large amount of cluster control link bandwidth.

A higher-bandwidth cluster control link helps the cluster to converge faster when there are membership changes and prevents throughput bottlenecks.

Note If your cluster has large amounts of asymmetric (rebalanced) traffic, then you should increase the cluster control link size.

Cluster Control Link Redundancy

The following diagram shows how to use an EtherChannel as a cluster control link in a Virtual Switching System (VSS), Virtual Port Channel (vPC), StackWise, or StackWise Virtual environment. All links in the EtherChannel are active. When the switch is part of a redundant system, then you can connect firewall interfaces within the same EtherChannel to separate switches in the redundant system. The switch interfaces are members of the same EtherChannel port-channel interface, because the separate switches act like a single switch. Note that this EtherChannel is device-local, not a Spanned EtherChannel.

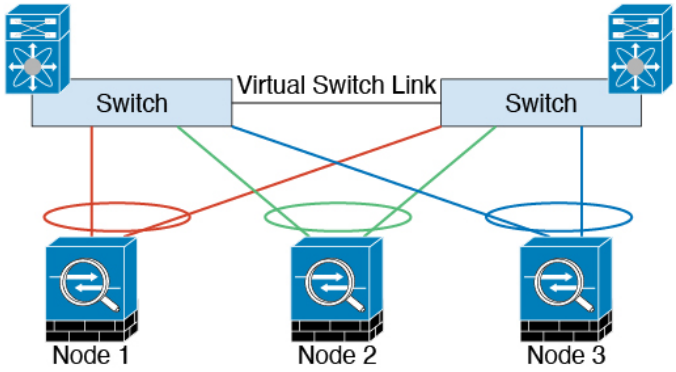

Cluster Control Link Reliability

To ensure cluster control link functionality, be sure the round-trip time (RTT) between units is less than 20 ms. This maximum latency enhances compatibility with cluster members installed at different geographical sites. To check your latency, perform a ping on the cluster control link between units.

The cluster control link must be reliable, with no out-of-order or dropped packets; for example, for inter-site deployment, you should use a dedicated link.

Spanned EtherChannels (Recommended)

You can group one or more interfaces per chassis into an EtherChannel that spans all chassis in the cluster. The EtherChannel aggregates the traffic across all the available active interfaces in the channel.

For regular firewall interfaces:A Spanned EtherChannel can be configured in both routed and transparent firewall modes. In routed mode, the EtherChannel is configured as a routed interface with a single IP address. In transparent mode, the IP address is assigned to the BVI, not to the bridge group member interface.

The EtherChannel inherently provides load balancing as part of basic operation.

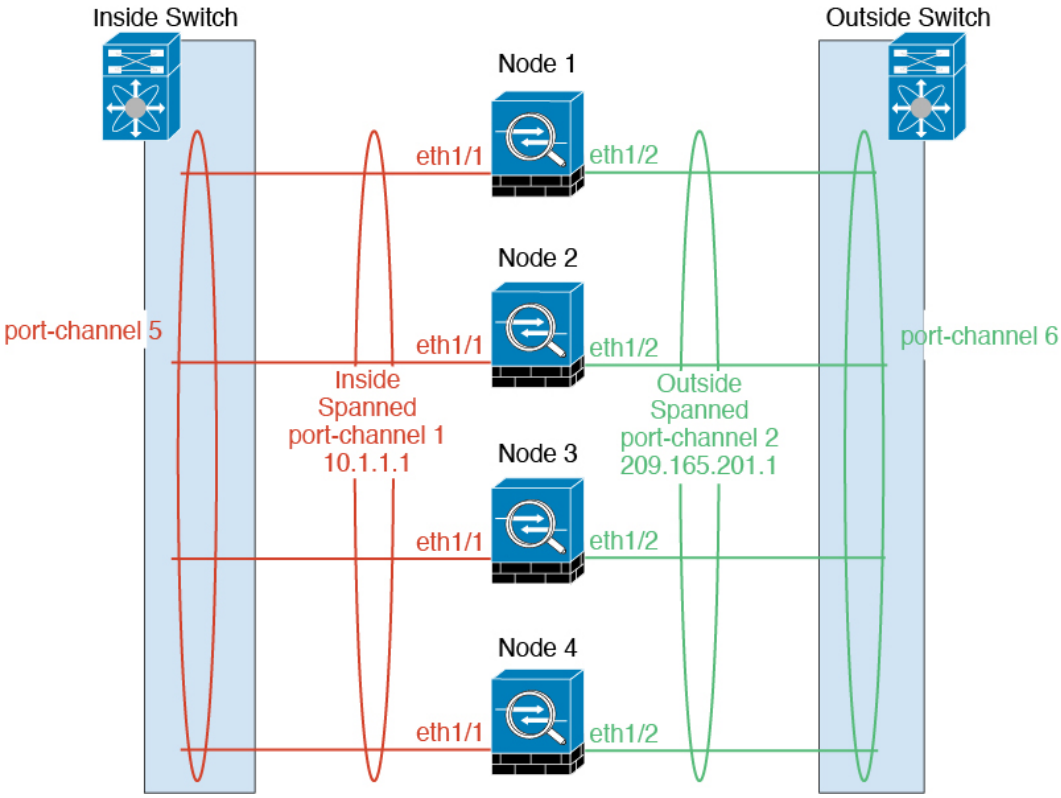

Spanned EtherChannel Benefits

The EtherChannel method of load-balancing is recommended over other methods for the following benefits:

- Faster failure discovery.
- Faster convergence time. Individual interfaces rely on routing protocols to load-balance traffic, and routing protocols often have slow convergence during a link failure.
- Ease of configuration.

Guidelines for Maximum Throughput

To achieve maximum throughput, we recommend the following:

- Use a load-balancing hash algorithm that is "symmetric," meaning that packets from both directions will have the same hash and will be sent to the same Firewall Threat Defense in the Spanned EtherChannel. We recommend using the source and destination IP address (the default) or the source and destination port as the hashing algorithm.
- Use the same type of line cards when connecting the Firewall Threat Defenses to the switch so that hashing algorithms applied to all packets are the same.

Load Balancing

The EtherChannel link is selected using a proprietary hash algorithm, based on source or destination IP addresses and TCP and UDP port numbers.

Note On the switch, we recommend that you use one of the following algorithms: **source-dest-ip** or **source-dest-ip-port** (see the Cisco Nexus OS or Cisco IOS **port-channel load-balance** command). Do not use a **vlan** keyword in the load-balance algorithm because it can cause unevenly distributed traffic to the nodes in a cluster.

The number of links in the EtherChannel affects load balancing.

Symmetric load balancing is not always possible. If you configure NAT, then forward and return packets will have different IP addresses and/or ports. Return traffic will be sent to a different unit based on the hash, and the cluster will have to redirect most returning traffic to the correct unit.

EtherChannel Redundancy

The EtherChannel has built-in redundancy. It monitors the line protocol status of all links. If one link fails, traffic is re-balanced between remaining links. If all links in the EtherChannel fail on a particular unit, but other units are still active, then the unit is removed from the cluster.

Connecting to a Redundant Switch System

You can include multiple interfaces per Firewall Threat Defense in the Spanned EtherChannel. Multiple interfaces per Firewall Threat Defense are especially useful for connecting to both switches in a VSS, vPC, StackWise, or StackWise Virtual system.

Depending on your switches, you can configure up to 32 active links in the spanned EtherChannel. This feature requires both switches in the vPC to support EtherChannels with 16 active links each (for example the Cisco Nexus 7000 with F2-Series 10 Gigabit Ethernet Module).

For switches that support 8 active links in the EtherChannel, you can configure up to 16 active links in the spanned EtherChannel when connecting to two switches in a redundant system.

The following figure shows a 16-active-link spanned EtherChannel in a 4-node cluster and an 8-node cluster.

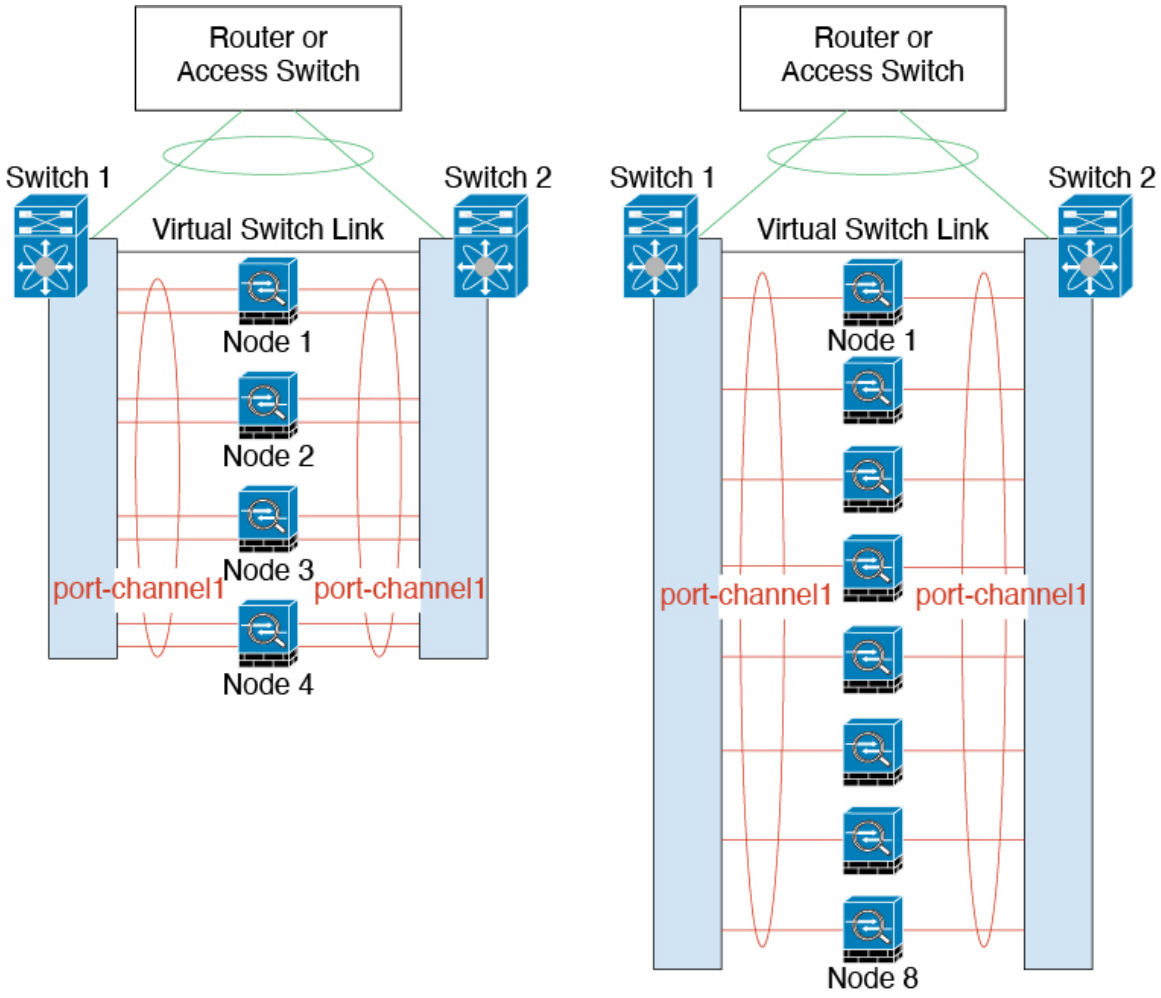

Individual Interfaces (Routed Firewall Mode Only)

Individual interfaces are normal routed interfaces, each with their own *Local IP address* used for routing. The *Main cluster IP address* for each interface is a fixed address that always belongs to the control node. When the control node changes, the Main cluster IP address moves to the new control node, so management of the cluster continues seamlessly.

IPS-only interfaces (inline sets and passive interfaces) are not supported as Individual interfaces.

Because interface configuration must be configured only on the control node, you configure a pool of IP addresses to be used for a given interface on the cluster nodes, including one for the control node.

Load balancing must be configured separately on the upstream switch.

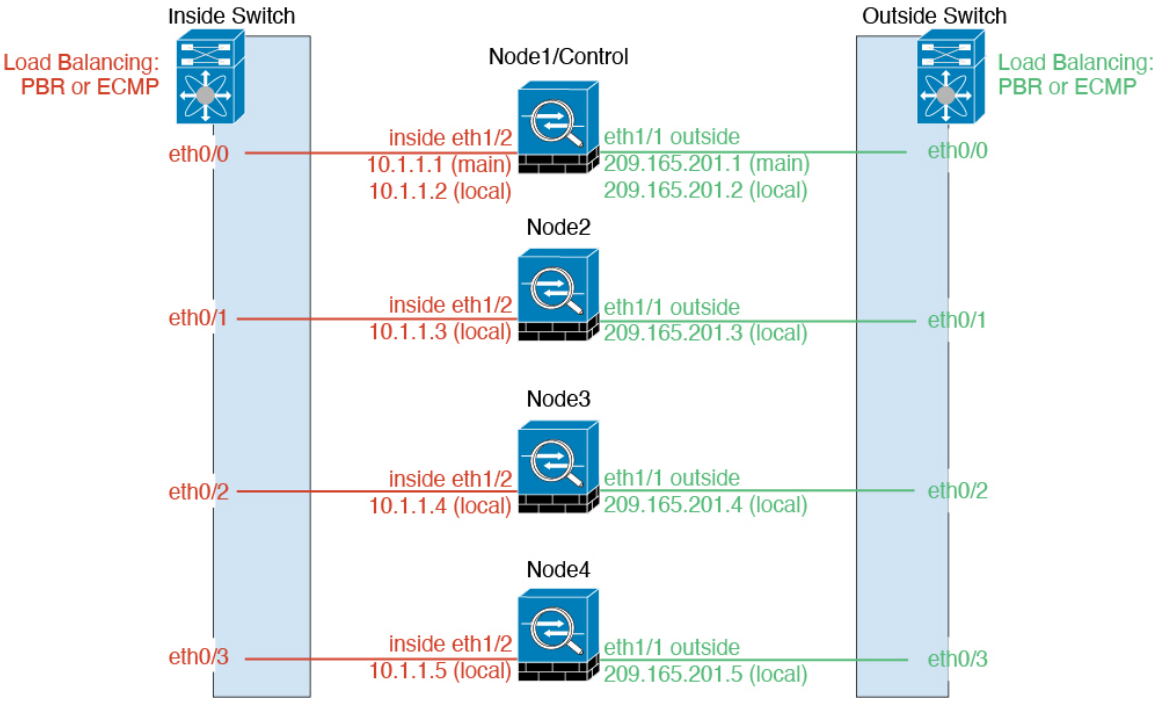

Policy-Based Routing

When using Individual interfaces, each Firewall Threat Defense interface maintains its own IP address and MAC address. One method of load balancing is Policy-Based Routing (PBR).

We recommend this method if you are already using PBR, and want to take advantage of your existing infrastructure.

PBR makes routing decisions based on a route map and ACL. You must manually divide traffic between all Firewall Threat Defenses in a cluster. Because PBR is static, it may not achieve the optimum load balancing result at all times. To achieve the best performance, we recommend that you configure the PBR policy so that forward and return packets of a connection are directed to the same Firewall Threat Defense. For example, if you have a Cisco router, redundancy can be achieved by using Cisco IOS PBR with Object Tracking. Cisco IOS Object Tracking monitors each Firewall Threat Defense using ICMP ping. PBR can then enable or disable route maps based on reachability of a particular Firewall Threat Defense. See the following URLs for more details:

http://www.cisco.com/c/en/us/solutions/data-center-virtualization/intelligent-traffic-director/index.html

http://www.cisco.com/en/US/products/ps6599/products_white_paper09186a00800a4409.shtml

Equal-Cost Multi-Path Routing

When using Individual interfaces, each Firewall Threat Defense interface maintains its own IP address and MAC address. One method of load balancing is Equal-Cost Multi-Path (ECMP) routing.

We recommend this method if you are already using ECMP, and want to take advantage of your existing infrastructure.

ECMP routing can forward packets over multiple “best paths” that tie for top place in the routing metric. Like EtherChannel, a hash of source and destination IP addresses and/or source and destination ports can be used to send a packet to one of the next hops. If you use static routes for ECMP routing, then the Firewall Threat

Defense failure can cause problems; the route continues to be used, and traffic to the failed Firewall Threat Defense will be lost. If you use static routes, be sure to use a static route monitoring feature such as Object Tracking. We recommend using dynamic routing protocols to add and remove routes, in which case, you must configure each Firewall Threat Defense to participate in dynamic routing.

Cisco Intelligent Traffic Director (Routed Firewall Mode Only)

When using Individual interfaces, each Firewall Threat Defense interface maintains its own IP address and MAC address. Intelligent Traffic Director (ITD) is a high-speed hardware load-balancing solution for Nexus 5000, 6000, 7000, and 9000 switch series. In addition to fully covering the functional capabilities of traditional PBR, it offers a simplified configuration workflow and multiple additional features for a more granular load distribution.

ITD supports IP stickiness, consistent hashing for bi-directional flow symmetry, virtual IP addressing, health monitoring, sophisticated failure handling policies with N+M redundancy, weighted load-balancing, and application IP SLA probes including DNS. Due to the dynamic nature of load-balancing, it achieves a more even traffic distribution across all cluster nodes as compared to PBR. In order to achieve bi-directional flow symmetry, we recommend configuring ITD such that forward and return packets of a connection are directed to the same Firewall Threat Defense. See the following URL for more details:

https://www.cisco.com/c/dam/en/us/td/docs/switches/datacenter/sw/design/itd_deployment/ITD_ASA_Deployment_Guide.pdf

Cable and Add Devices to the Firewall Management Center

Before configuring clustering, you need to prepare your devices. In particular, the cluster will not come up unless all nodes can communicate over the cluster control link. Therefore, before you form the cluster, the cluster control link must be ready to go.

Procedure

Step 1 Cable the cluster control link network, management network, and data networks.

Step 2 Configure the upstream and downstream equipment.

a) For the cluster control link network, set the MTU to be at least 100 bytes higher than the data interface MTU.

By default, the data interface MTU is 1500 bytes, so the cluster control link MTU on the cluster node will be set to 1600 bytes. If you use higher MTUs on your data interfaces, increase the cluster control link MTU on connecting switches accordingly.

b) Configure cluster control link interfaces on upstream and downstream equipment, including for an optional EtherChannel.

See Cluster Control Link Interfaces and Network, on page 465 for cluster control link requirements.

c) Configure data interfaces on upstream and downstream equipment, including Spanned EtherChannels, if you choose that cluster interface mode.

See About Cluster Interfaces, on page 464 for information about how to cable Spanned EtherChannels.

Step 3 Add each node to the Firewall Management Center as a standalone device in the same domain and group.

See Add a Device Using a Registration Key—Basic Configuration, on page 31. You can create a cluster with a single device, and then add more nodes later. The initial settings (licensing, access control policy) that you set when you add a device will be inherited by all cluster nodes from the control node. You will choose the control node when forming the cluster.

Step 4 Enable the cluster control link on the device you want to be the control node.

When you add the other nodes, they will inherit the cluster control link configuration.

Note
Do *not* configure the name or IP addressing for the cluster control link. The MTU of the cluster control link interface is automatically set to 100 bytes more than the highest data interface MTU when you form the cluster, so you do not need to set it now. However, we do not recommend setting the cluster control link MTU between 2561 and 8362; due to block pool handling, this MTU size is not optimal for system operation. If the MTU is set in this range when you add the cluster, we recommend returning to the **Interfaces** page and manually increasing it above 8362. When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the ping fails, a notification is generated so you can fix the MTU mismatch on connecting switches and try again.

a) On the device you want to be the control node, choose **Devices** > **Device Management**, and click **Edit** (✎).
b) Click **Interfaces**.
c) Enable the interface. If you are going to use an EtherChannel for the cluster control link, enable all members. See Enable the Physical Interface and Configure Ethernet Settings, on page 778.

Figure 197: Enable the Cluster Control Link Interface

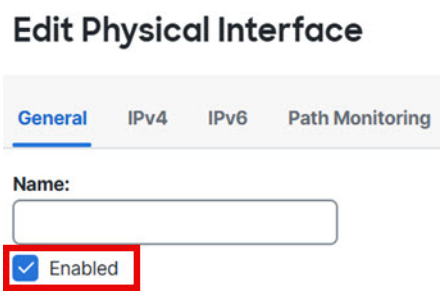

d) (Optional) Add an EtherChannel. See Configure an EtherChannel, on page 786.

We recommend using the On mode for cluster control link member interfaces to reduce unnecessary traffic on the cluster control link (Active mode is the default). The cluster control link does not need the overhead of LACP traffic because it is an isolated, stable network. **Note:** We recommend setting *data* EtherChannels to Active mode.

e) Click **Save** and then **Deploy** to deploy the interface changes to the control node.

Create a Cluster

Form a cluster from one or more devices in the Firewall Management Center.

Procedure

Step 1 Choose **Devices** > **Device Management**, and then choose **Add** > **Cluster**.

The **Add Cluster Wizard** appears.

Figure 198: Add Cluster Wizard

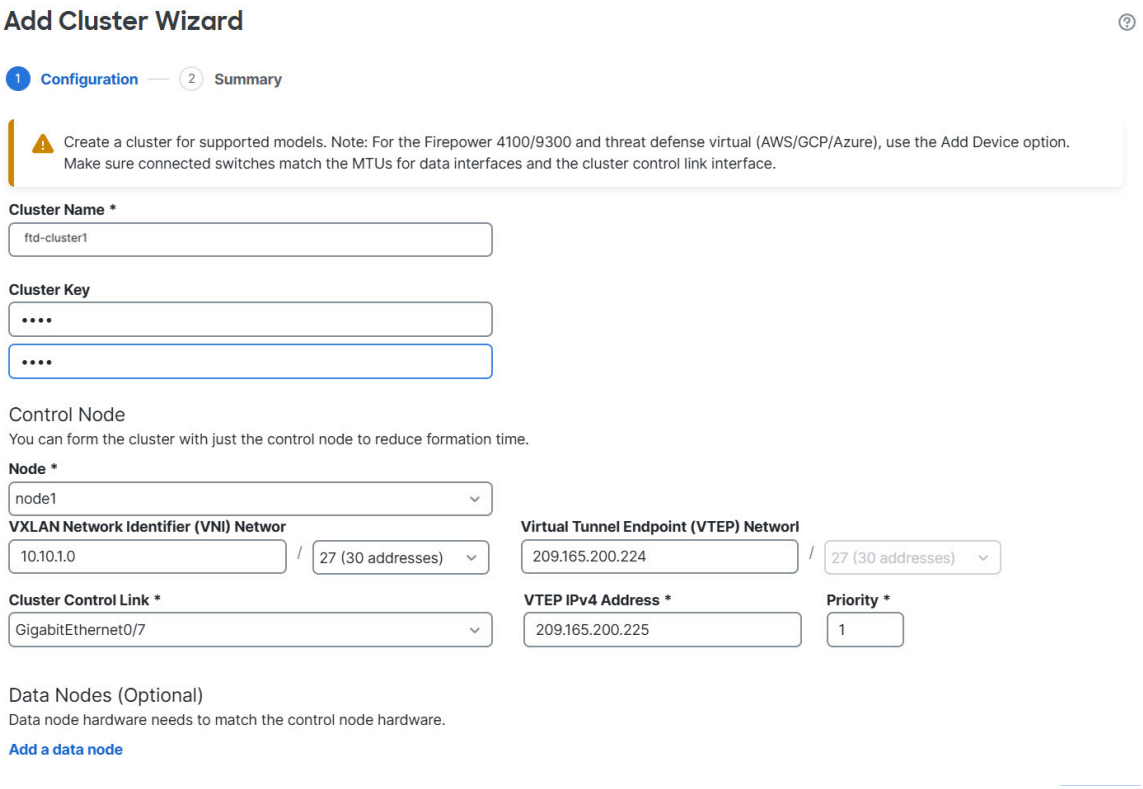

Step 2 Specify a **Cluster Name** and an authentication **Cluster Key** for control traffic.

- **Cluster Name**—An ASCII string from 1 to 38 characters.
- **Cluster Key**—An ASCII string from 1 to 63 characters. The **Cluster Key** value is used to generate the encryption key. This encryption does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear.

Step 3 For the **Control Node**, set the following:

- **Node**—Choose the device that you want to be the control node initially. When the Firewall Management Center forms the cluster, it will add this node to the cluster first so it will be the control node.

  Note

  If you see an **Error** ( ) icon next to the node name, click the icon to view configuration issues. You must cancel cluster formation, resolve the issues, and then return to cluster formation. For example:

Figure 199: Configuration Issues

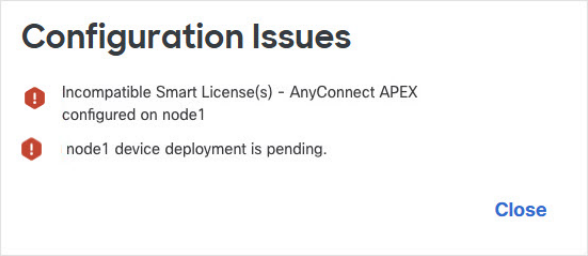

To resolve the above issues, remove the unsupported VPN license and deploy pending configuration changes to the device.

- **Cluster Control Link Network**—Specify an IPv4 subnet; IPv6 is not supported for this interface. Specify a **24**, **25**, **26**, or **27** subnet.
- **Cluster Control Link**—Choose the physical interface or EtherChannel you want to use for the cluster control link.

  Note
  The MTU of the cluster control link interface is automatically set to 100 bytes more than the highest data interface MTU; by default, the MTU is 1600 bytes. We do not recommend setting the cluster control link MTU between 2561 and 8362; due to block pool handling, this MTU size is not optimal for system operation. If the MTU is set in this range when you add the cluster, we recommend increasing the MTU above 8362 on the **Devices** > **Device Management** > **Interfaces** page.

  Make sure you configure switches connected to the cluster control link to the correct (higher) MTU; otherwise, cluster formation will fail. When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the ping fails, a notification is generated so you can fix the MTU mismatch on connecting switches and try again.

- **Cluster Control Link IPv4 Address**—This field will be auto-populated with the first address on the cluster control link network. You can edit the host address if desired.
- **Priority**—Set the priority of this node for control node elections. The priority is between 1 and 100, where 1 is the highest priority. Even if you set the priority to be lower than other nodes, this node will still be the control node when the cluster is first formed.
- **Site ID**—(FlexConfig feature) Enter the site ID for this node between 1 and 8. A value of 0 disables inter-site clustering. Additional inter-site cluster customizations to enhance redundancy and stability, such as director localization, site redundancy, and cluster flow mobility, are only configurable using the FlexConfig feature.

Step 4 For the **Cluster Mode**, choose **Spanned EtherChannel Mode** or **Individual Interface Mode**.

Step 5 For **Data Nodes (Optional)**, click **Add a data node** to add a node to the cluster.

You can form the cluster with only the control node for faster cluster formation, or you can add all nodes now. Set the following for each data node:

- **Node**—Choose the device that you want to add.

  Note

If you see an **Error** ( ) icon next to the node name, click the icon to view configuration issues. You must cancel cluster formation, resolve the issues, and then return to cluster formation.

- **Cluster Control Link IPv4 Address**—This field will be auto-populated with the next address on the cluster control link network. You can edit the host address if desired.
- **Priority**—Set the priority of this node for control node elections. The priority is between 1 and 100, where 1 is the highest priority.
- **Site ID**—(FlexConfig feature) Enter the site ID for this node between 1 and 8. A value of 0 disables inter-site clustering. Additional inter-site cluster customizations to enhance redundancy and stability, such as director localization, site redundancy, and cluster flow mobility, are only configurable using the FlexConfig feature.

Step 6 Click **Continue**. Review the **Summary**, and then click **Save**.

The cluster name shows on the **Devices** > **Device Management** page; expand the cluster to see the cluster nodes.

Figure 200: Cluster Management

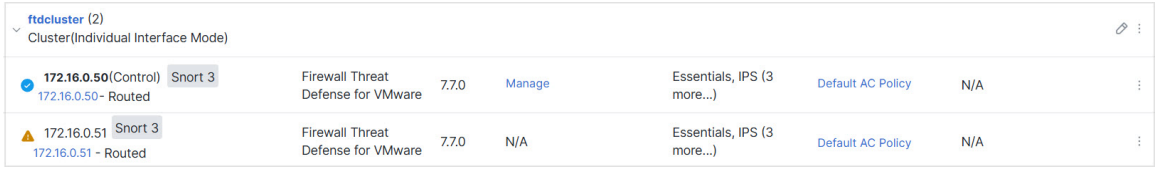

A node that is currently registering shows the loading icon.

Figure 201: Node Registration

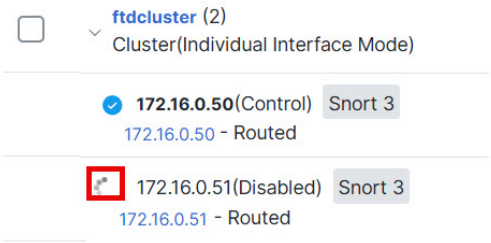

You can monitor cluster node registration by clicking the **Notifications** icon and choosing **Tasks**. The Firewall Management Center updates the Cluster Registration task as each node registers.

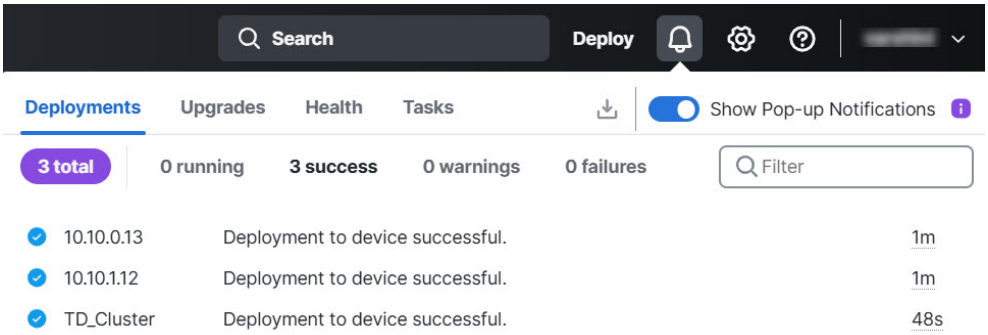

Step 7 Configure device-specific settings by clicking the **Edit** (✎) for the cluster.

Most configuration can be applied to the cluster as a whole, and not nodes in the cluster. For example, you can change the display name per node, but you can only configure interfaces for the whole cluster.

Step 8 On the **Devices** > **Device Management** > **Cluster** screen, you see **General** and other settings for the cluster.

Figure 202: Cluster Settings

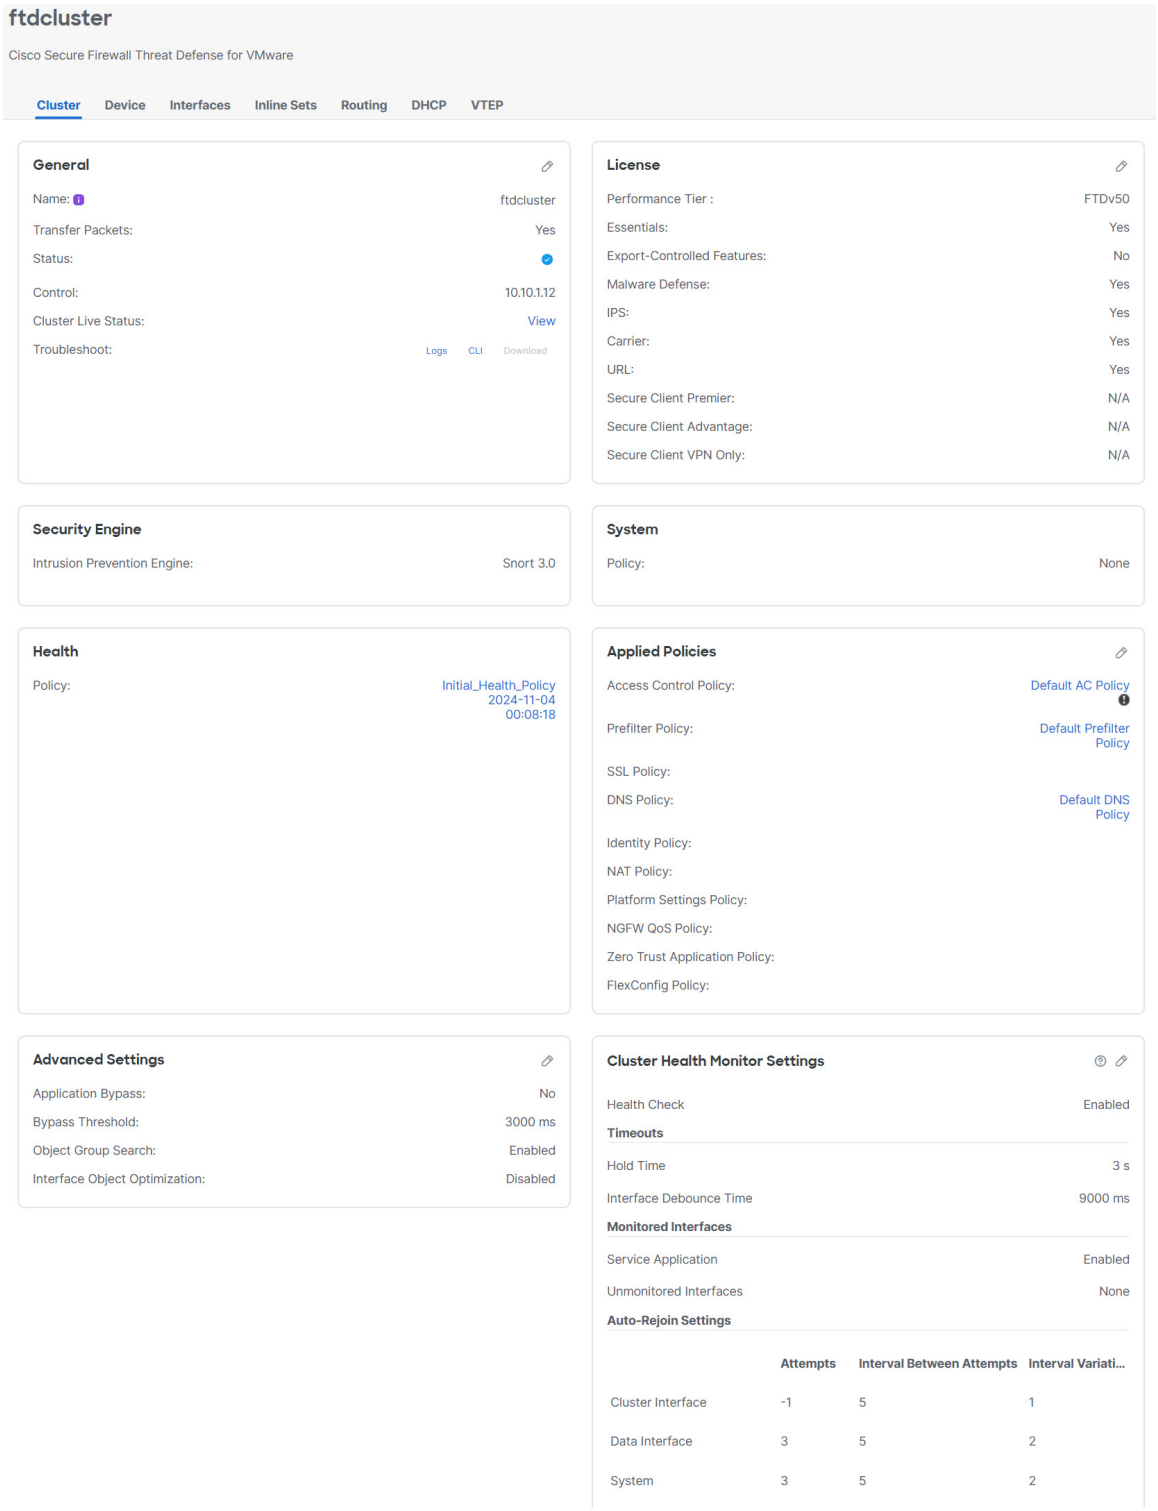

See the following cluster-specific items in the **General** area:

- **General > Name**—Change the cluster display name by clicking the **Edit** ( ).

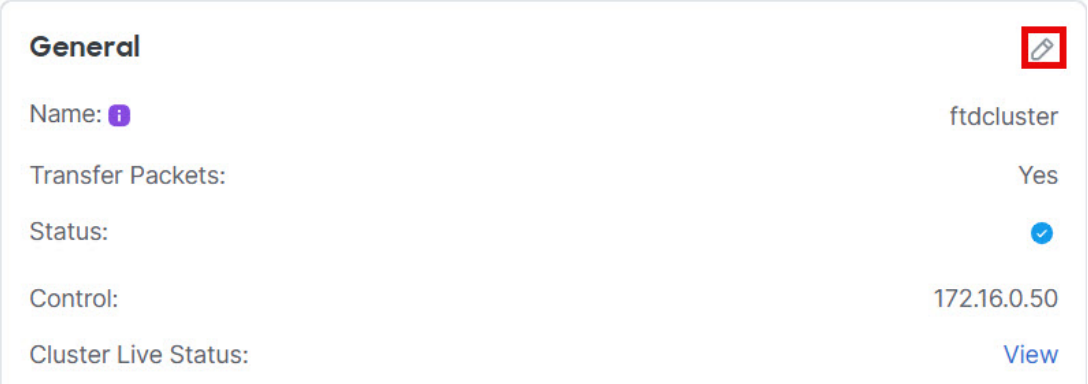

Then set the **Name** field.

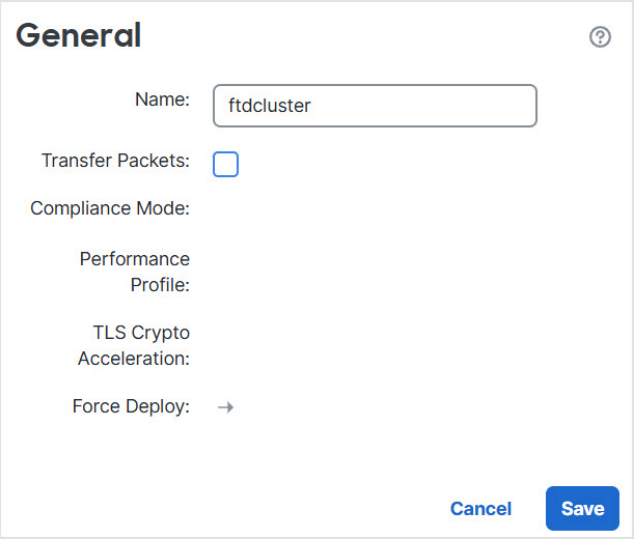

- **General > Cluster Live Status**—Click the **View** link to open the **Cluster Status** dialog box.

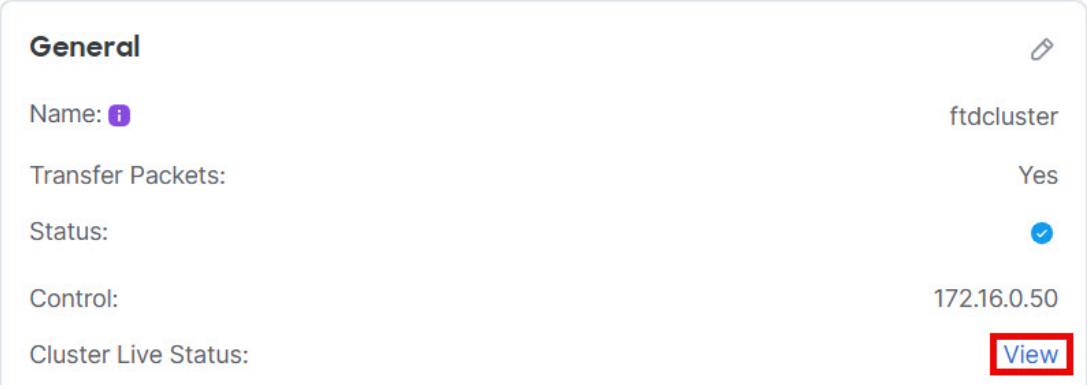

The **Cluster Status** dialog box also lets you retry data unit registration by clicking **Reconcile All**. You can also ping the cluster control link from a node. See Perform a Ping on the Cluster Control Link, on page 505.

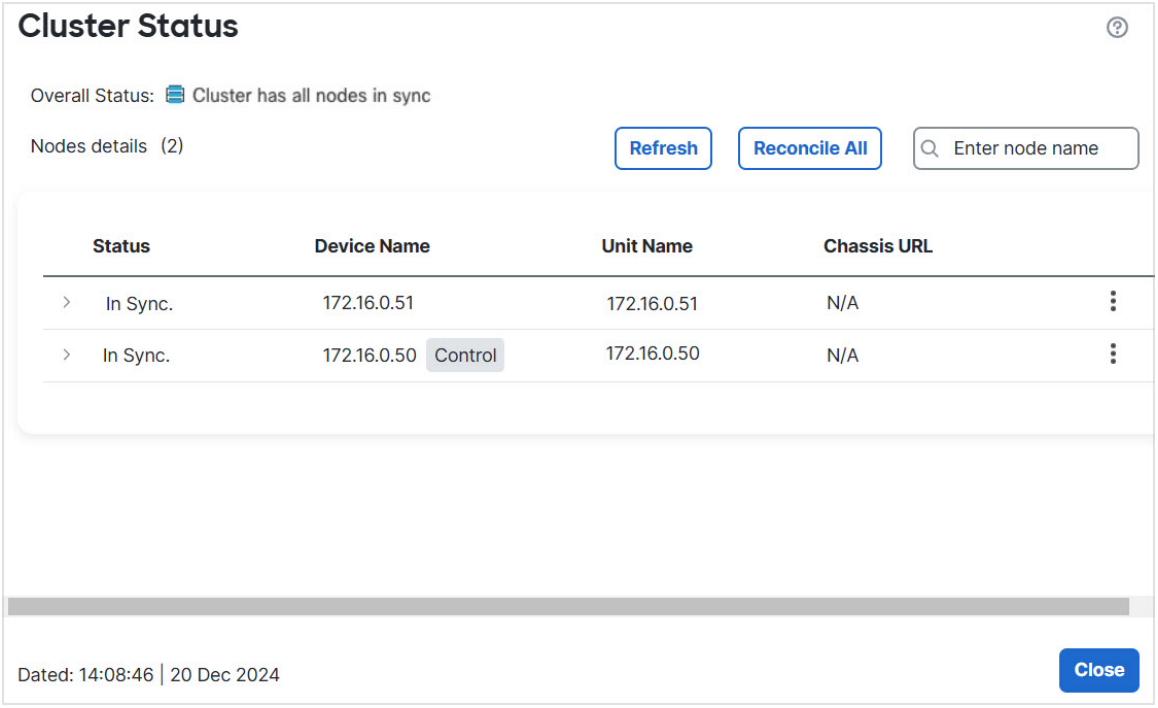

- **General > Troubleshoot**—You can generate and download troubleshooting logs, and you can view cluster CLIs. See Troubleshooting the Cluster, on page 505.

Figure 203: Troubleshoot

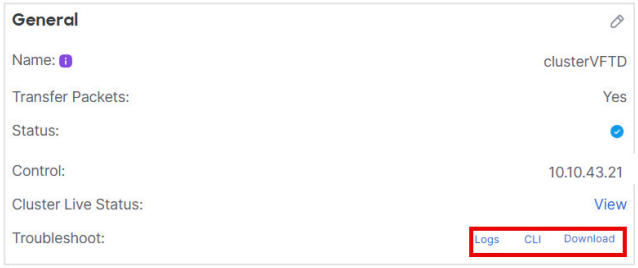

Step 9 On the **Devices** > **Device Management** > **Devices**, you can choose each member in the cluster from the top right drop-down menu and configure the following settings.

Figure 204: Device Settings

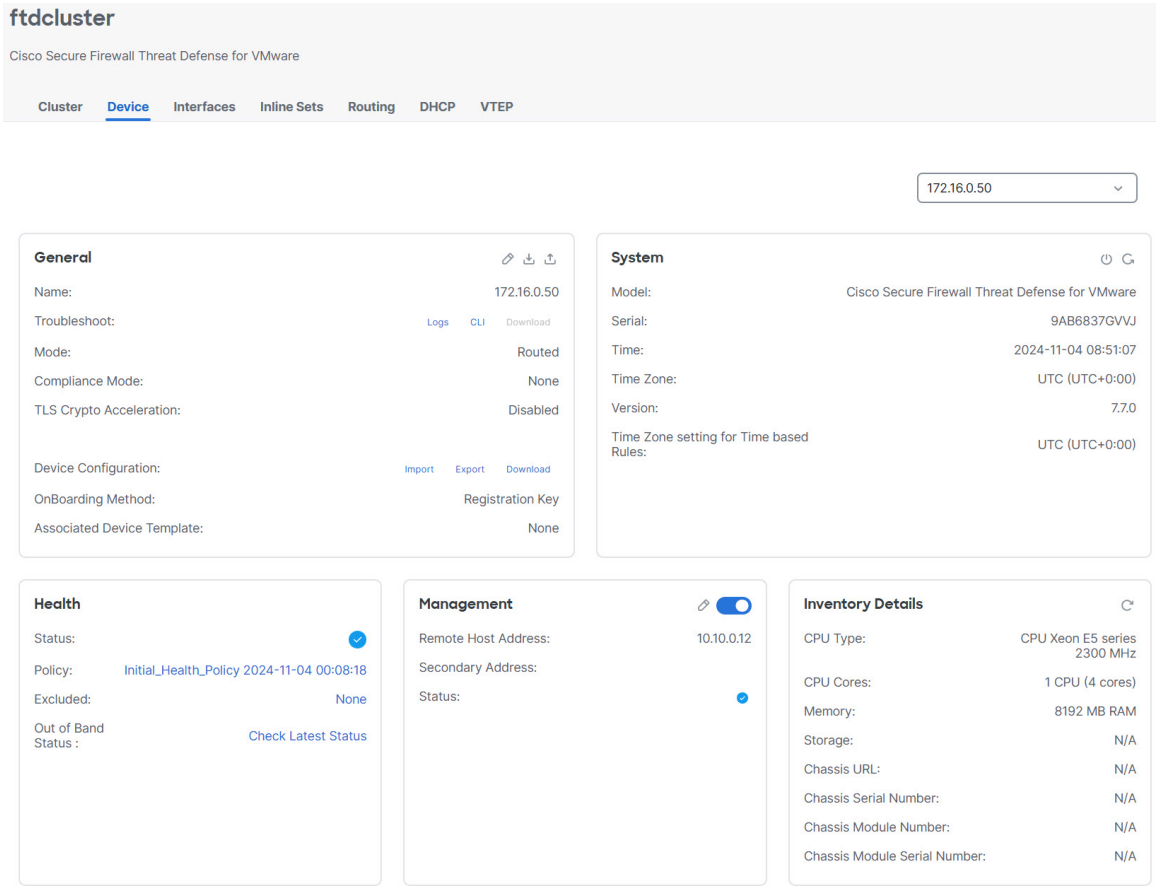

Figure 205: Choose Node

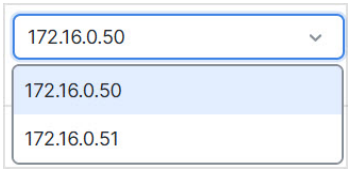

- **General > Name**—Change the cluster member display name by clicking the **Edit** (✎).

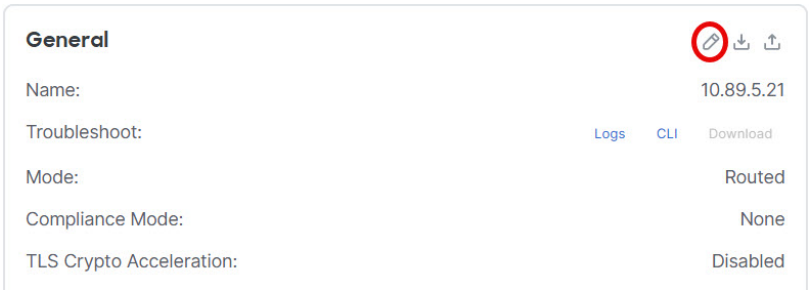

Then set the **Name** field.

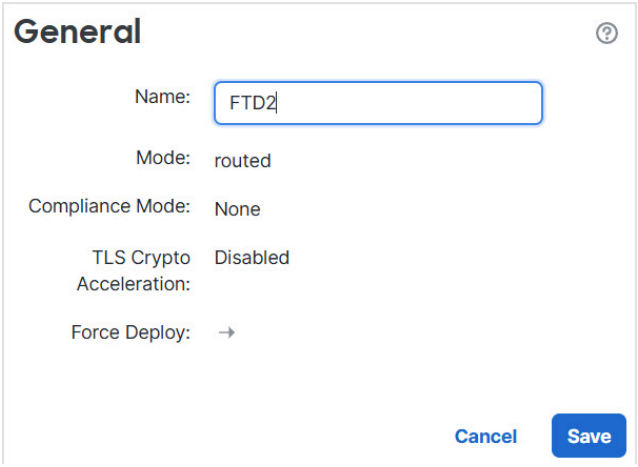

- **Management > Host**—If you change the management IP address in the device configuration, you must match the new address in the Firewall Management Center so that it can reach the device on the network. First disable the connection, edit the **Host** address in the **Management** area, then re-enable the connection.

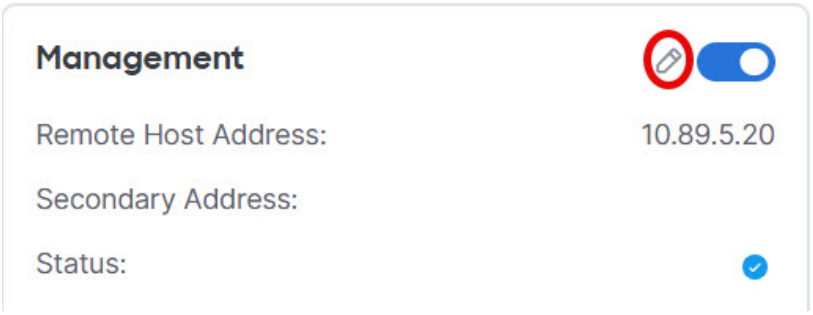

Configure Interfaces

You can configure data interfaces as either Spanned EtherChannels or as Individual interfaces. Each method uses a different load-balancing mechanism. You cannot configure both types in the same configuration.

Configure Spanned EtherChannels

Configure data interfaces as Spanned EtherChannels.

Procedure

Step 1 Choose **Devices** > **Device Management**, and click **Edit** (✎) next to the cluster.

Step 2 Click **Interfaces**.

Step 3 Configure Spanned EtherChannel data interfaces.

a) Configure one or more EtherChannels. See Configure an EtherChannel, on page 786.

You can include one or more member interfaces in the EtherChannel. Because this EtherChannel is spanned across all of the nodes, you only need one member interface per node; however, for greater throughput and redundancy, multiple members are recommended.

b) (Optional) For regular firewall interfaces, configure VLAN subinterfaces on the EtherChannel. The rest of this procedure applies to the subinterfaces. See Add a Subinterface, on page 836.

c) Click **Edit** (✎) for the EtherChannel interface.

d) Configure the name and other parameters. For regular firewall interfaces, see Configure Routed Mode Interfaces, on page 856 or, for transparent mode, Configure Bridge Group Interfaces, on page 861. For IPS-only interfaces, see Inline Sets and Passive Interfaces, on page 893.

- If the cluster control link interface MTU is not at least 100 bytes higher than the data interface MTU, you will see an error that you must reduce the MTU of the data interface. By default, the cluster control link MTU is 1600 bytes. If you want to increase the MTU of data interfaces, first increase the cluster control link MTU. Note that we do not recommend setting the cluster control link MTU between 2561 and 8362; due to block pool handling, this MTU size is not optimal for system operation.
- For routed mode, DHCP, PPPoE, IPv6 autoconfig and manual link-local addresses are not supported. For point-to-point connections, you can specify a 31-bit subnet mask (255.255.255.254). In this case, no IP addresses are reserved for the network or broadcast addresses.

e) Set a unique, manual global MAC address for the EtherChannel. Click **Advanced**, and in the **Active Mac Address** field, enter a MAC address in H.H.H format, where H is a 16-bit hexadecimal digit.

For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE. The MAC address must not have the multicast bit set, that is, the second hexadecimal digit from the left cannot be an odd number.

Do not set the **Standby Mac Address**; it is ignored.

You must configure a unique MAC address not currently in use on your network for a Spanned EtherChannel to avoid potential network connectivity problems. With a manually-configured MAC address, the MAC address stays with the current control unit. If you do not configure a MAC address, then if the control unit changes, the new control unit uses a new MAC address for the interface, which can cause a temporary network outage.

f) Click **OK**. Repeat the above steps for other data interfaces.

Step 4 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Individual Interfaces

Individual interfaces are normal routed interfaces, each with their own IP address taken from a pool of IP addresses. The Main cluster IP address is a fixed address for the cluster that always belongs to the control node.

Individual management interfaces let you SSH directly to each unit if necessary, while a Spanned EtherChannel interface only allows connection to the control node.

IPS-only interfaces (inline sets or passive interfaces) are not supported for Individual interfaces.

Before you begin

- You must be in Individual interface mode.
- Individual interfaces require you to configure load balancing on neighbor devices. External load balancing is not required for the management interface.
- (Optional) Configure the interface as a device-local EtherChannel interface, and/or configure subinterfaces.

  For an EtherChannel, this EtherChannel is local to the unit, and is not a Spanned EtherChannel.

Procedure

Step 1 Choose **Objects** > **Object Management** > **Address Pools** to add an IPv4 and/or IPv6 address pool. See Address Pools, on page 1374.

Include at least as many addresses as there are units in the cluster. The main IP address is not a part of this pool, but needs to be on the same network. You cannot determine the exact Local address assigned to each unit in advance.

Figure 206: Add Address Pool

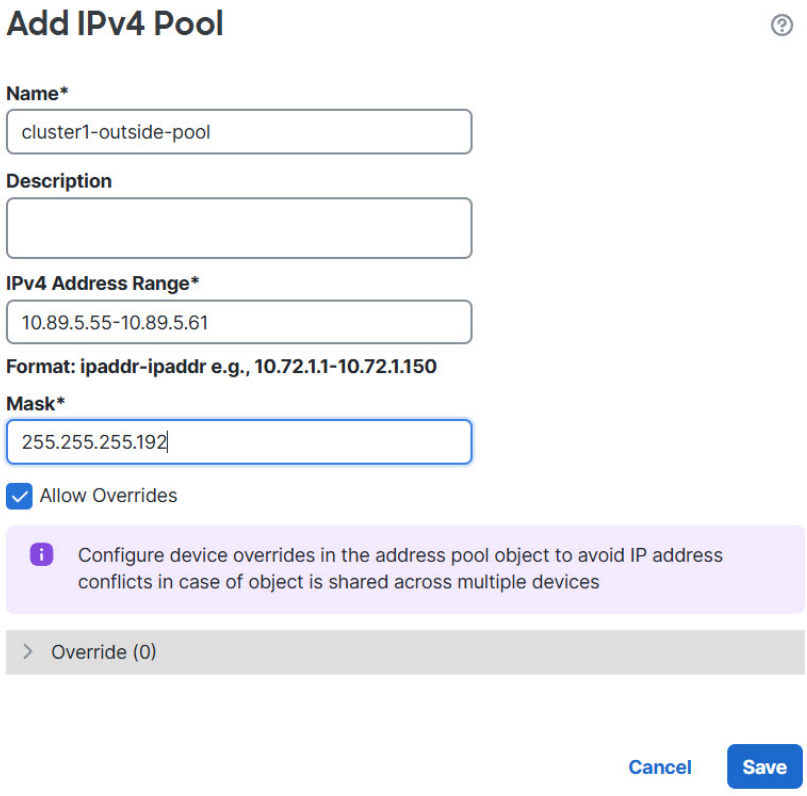

Note
Although not common, if you want to set the MAC addresses manually, you can also add a MAC address pool object.

Step 2 On **Devices** > **Device Management** > **Interfaces**, click **Edit** (✎) for the interface you want to configure.

Step 3 On the **IPv4** page, enter the **Virtual IP Address** and mask. This main ("virtual") IP address is a fixed address for the cluster, and always belongs to the control node.

DHCP and PPPoE are not supported. For point-to-point connections, you can specify a 31-bit subnet mask (255.255.255.254). In this case, no IP addresses are reserved for the network or broadcast addresses.

Figure 207: IPv4 Page

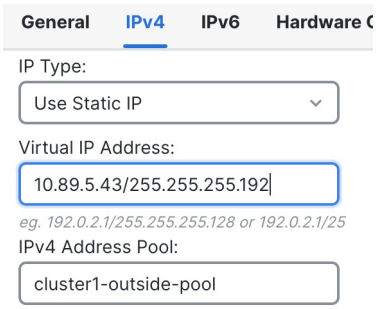

Step 4 From the **IPv4 Address Pool** drop-down list, choose the address pool you created.

Step 5 On **IPv6 > Basic**, from the **IPv6 Address Pool** drop-down list, choose the address pool you created.

IPv6 autoconfig and manual link-local addresses are not supported.

Step 6 Configure other interface settings as normal.

To set the MAC addresses manually, you can select the MAC address pool from the interface's **Advanced** page.

Note
If the cluster control link interface MTU is not at least 100 bytes higher than the data interface MTU, you will see an error that you must reduce the MTU of the data interface. By default, the cluster control link MTU is 1600 bytes. If you want to increase the MTU of data interfaces, first increase the cluster control link MTU. Note that we do not recommend setting the cluster control link MTU between 2561 and 8362; due to block pool handling, this MTU size is not optimal for system operation.

Configure Cluster Health Monitor Settings

The **Cluster Health Monitor Settings** section of the **Cluster** page displays the settings described in the table below.

Figure 208: Cluster Health Monitor Settings

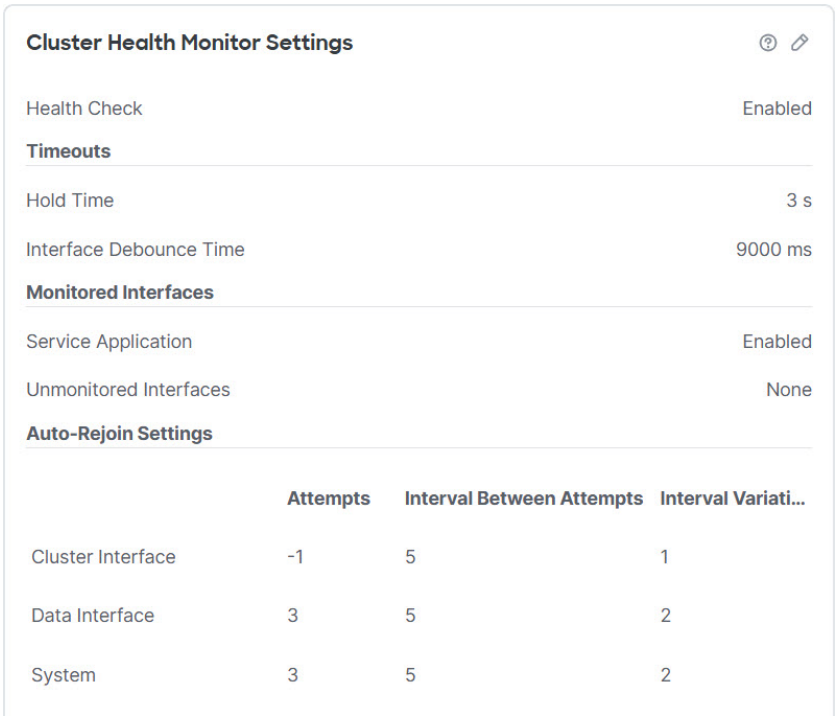

Table 31: Cluster Health Monitor Settings Section Table Fields

Field	Description
Timeouts	
Hold Time	Between .3 and 45 seconds; The default is 3 seconds. To determine node system health, the cluster nodes send heartbeat messages on the cluster control link to other nodes. If a node does not receive any heartbeat messages from a peer node within the hold time period, the peer node is considered unresponsive or dead.
Interface Debounce Time	Between 300 and 9000 ms. The default is 500 ms. The interface debounce time is the amount of time before the node considers an interface to be failed, and the node is removed from the cluster.
Monitored Interfaces	The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The amount of time before the node removes a member from the cluster depends on the type of interface and whether the node is an established node or is joining the cluster.
Service Application	Shows whether the Snort and disk-full processes are monitored.
Unmonitored Interfaces	Shows unmonitored interfaces.
Auto-Rejoin Settings	

Field	Description
Cluster Interface	Shows the auto-rejoin settings after a cluster control link failure.
Attempts	Between -1 and 65535. The default is -1 (unlimited). Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 1x the interval duration. Defines if the interval duration increases at each attempt.
Data Interfaces	Shows the auto-rejoin settings after a data interface failure.
Attempts	Between -1 and 65535. The default is 3. Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 2x the interval duration. Defines if the interval duration increases at each attempt.
System	Shows the auto-rejoin settings after internal errors. Internal failures include: application sync timeout; inconsistent application statuses; and so on.
Attempts	Between -1 and 65535. The default is 3. Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 2x the interval duration. Defines if the interval duration increases at each attempt.

Note If you disable the system health check, fields that do not apply when the system health check is disabled will not show.

You can change these settings from this section.

You can monitor any port-channel ID, single physical interface ID, as well as the Snort and disk-full processes. Health monitoring is not performed on VLAN subinterfaces or virtual interfaces such as VNIs or BVIs. You cannot configure monitoring for the cluster control link; it is always monitored.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the cluster you want to modify, click **Edit** ( ).

Step 3 Click **Cluster**.

Step 4 In the **Cluster Health Monitor Settings** section, click **Edit** ( ).

Step 5 Disable the system health check by clicking the **Health Check** slider .

Figure 209: Disable the System Health Check

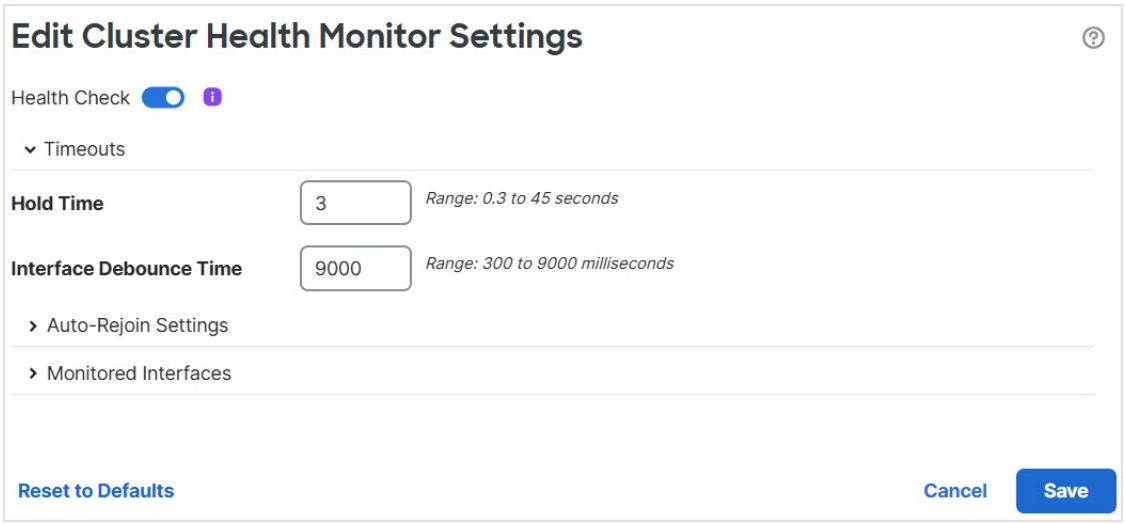

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 6 Configure the hold time and interface debounce time.

- **Hold Time**—Set the hold time to determine the amount of time between node heartbeat status messages, between .3 and 45 seconds; The default is 3 seconds.
- **Interface Debounce Time**—Set the debounce time between 300 and 9000 ms. The default is 500 ms. Lower values allow for faster detection of interface failures. Note that configuring a lower debounce time increases the chances of false-positives. When an interface status update occurs, the node waits the number of milliseconds specified before marking the interface as failed, and the node is removed from the cluster. In the case of an EtherChannel that transitions from a down state to an up state (for example, the switch reloaded, or the switch enabled an EtherChannel), a longer debounce time can prevent the interface from appearing to be failed on a cluster node just because another cluster node was faster at bundling the ports.

Step 7 Customize the auto-rejoin cluster settings after a health check failure.

Figure 210: Configure Auto-Rejoin Settings

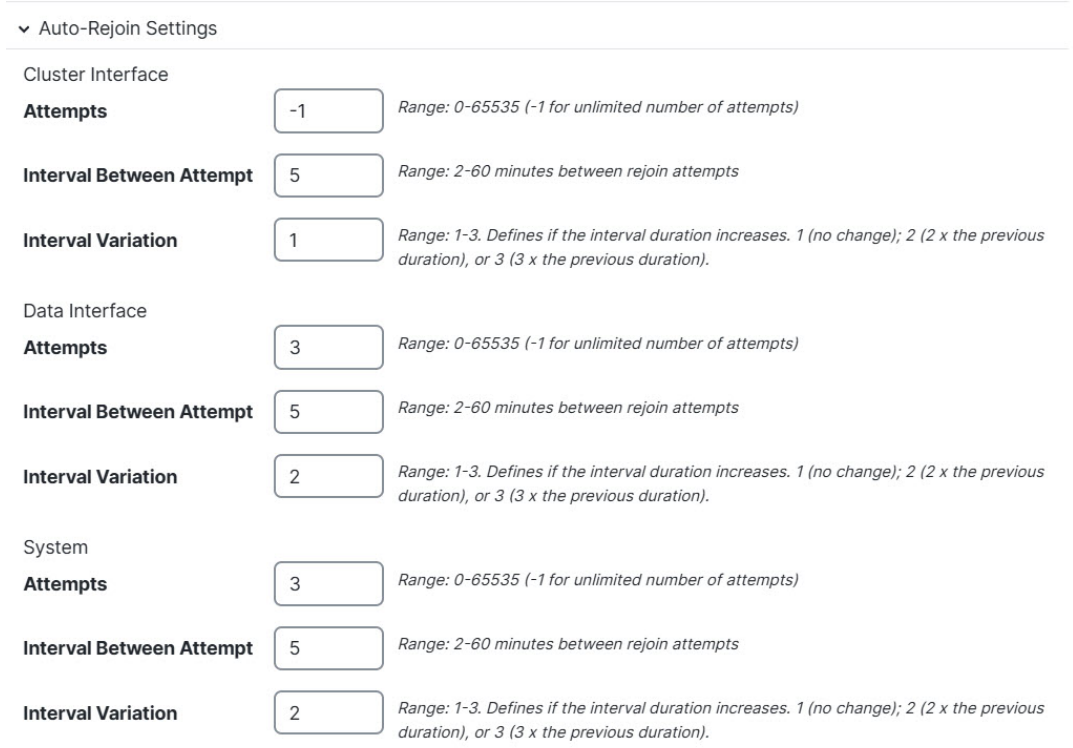

Set the following values for the **Cluster Interface**, **Data Interface**, and **System** (internal failures include: application sync timeout; inconsistent application statuses; and so on):

- **Attempts**—Sets the number of rejoin attempts, between -1 and 65535. **0** disables auto-rejoining. The default for the **Cluster Interface** is -1 (unlimited). The default for the **Data Interface** and **System** is 3.
- **Interval Between Attempts**—Defines the interval duration in minutes between rejoin attempts, between 2 and 60. The default value is 5 minutes. The maximum total time that the node attempts to rejoin the cluster is limited to 14400 minutes (10 days) from the time of last failure.
- **Interval Variation**—Defines if the interval duration increases. Set the value between 1 and 3: **1** (no change); **2** (2 x the previous duration), or **3** (3 x the previous duration). For example, if you set the interval duration to 5 minutes, and set the variation to 2, then the first attempt is after 5 minutes; the 2nd attempt is 10 minutes (2 x 5); the 3rd attempt 20 minutes (2 x 10), and so on. The default value is **1** for the **Cluster Interface** and **2** for the **Data Interface** and **System**.

Step 8 Configure monitored interfaces by moving interfaces in the **Monitored Interfaces** or **Unmonitored Interfaces** window. You can also check or uncheck **Enable Service Application Monitoring** to enable or disable monitoring of the Snort and disk-full processes.

Figure 211: Configure Monitored Interfaces

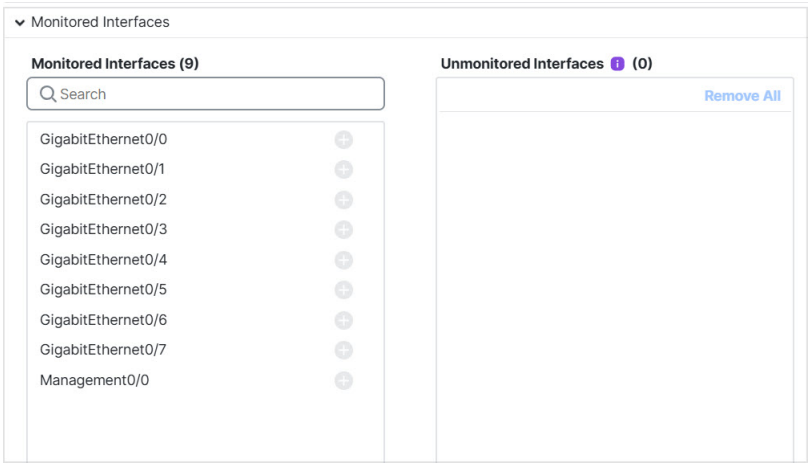

The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The amount of time before the node removes a member from the cluster depends on the type of interface and whether the node is an established node or is joining the cluster. Health check is enabled by default for all interfaces and for the Snort and disk-full processes.

You might want to disable health monitoring of non-essential interfaces.

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 9 Click **Save**.

Step 10 Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Manage Cluster Nodes

After you deploy the cluster, you can change the configuration and manage cluster nodes.

Add a New Cluster Node

You can add one or more new cluster nodes to an existing cluster.

Procedure

Step 1 Choose **Devices** > **Device Management**, click **More** (⋮) for the cluster, and choose **Add Nodes**.

Figure 212: Add Nodes

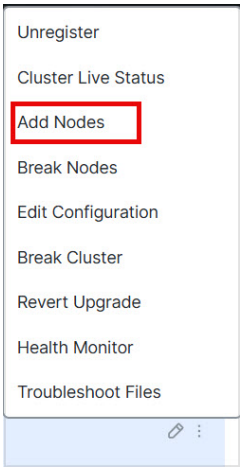

The **Manage Cluster Wizard** appears.

Step 2 From the **Node** menu, choose a device, adjust the IP address, priority, and Site ID if desired.

Figure 213: Manage Cluster Wizard

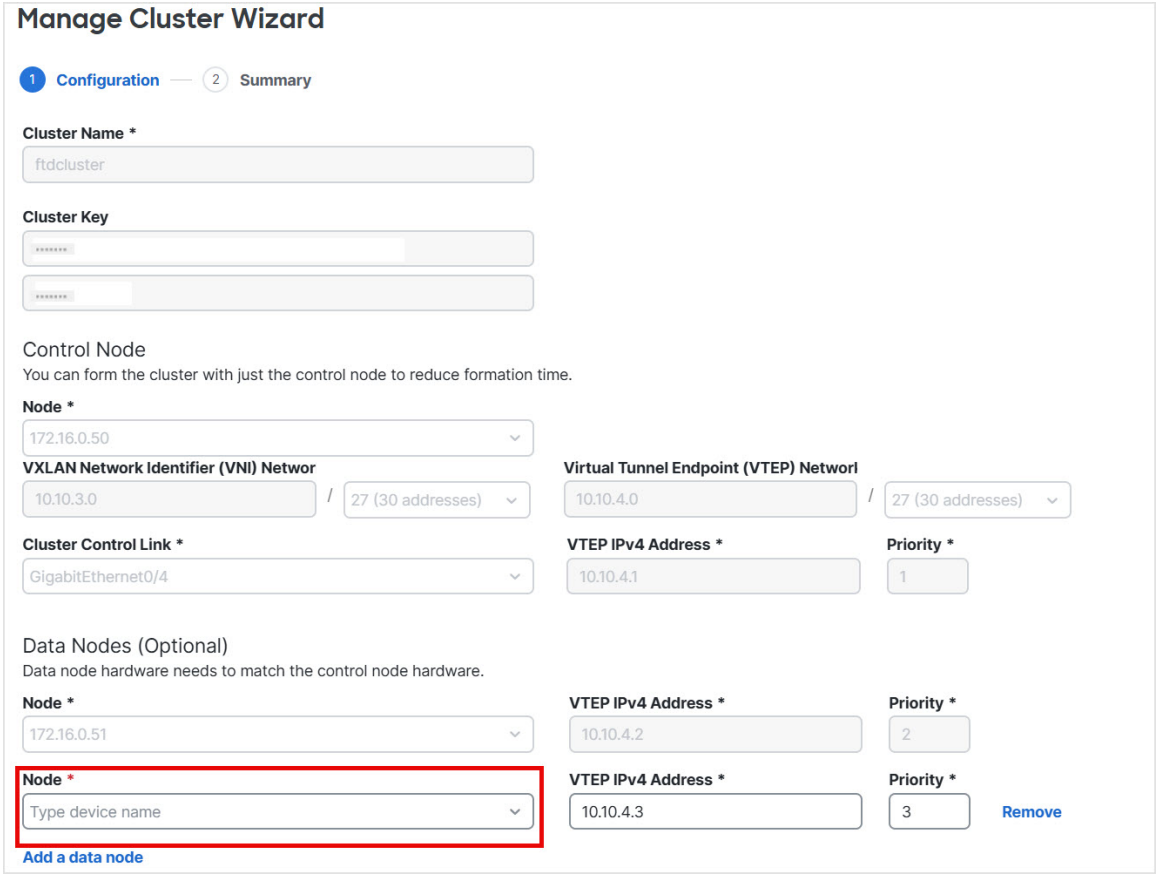

Step 3 To add additional nodes, click **Add a data node**.

Step 4 Click **Continue**. Review the **Summary**, and then click **Save**

The node that is currently registering shows the loading icon.

Figure 214: Node Registration

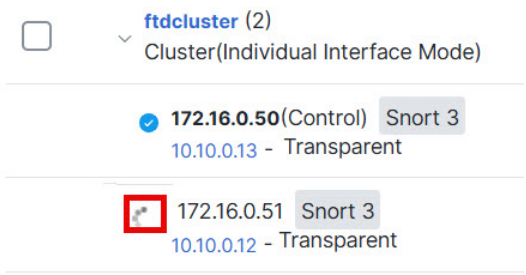

You can monitor cluster node registration by clicking the **Notifications** icon and choosing **Tasks**.

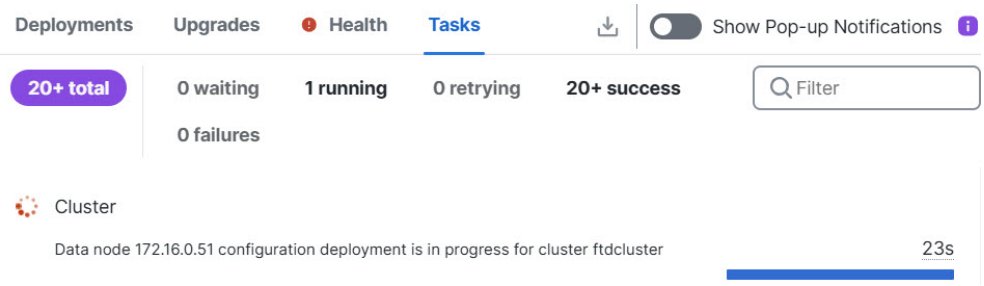

Break a Node

You can remove a node from the cluster so that it becomes a standalone device. You cannot break the control node unless you break the entire cluster. The data node has its configuration erased.

Procedure

Step 1 Choose **Devices** > **Device Management**, click **More** (⋮) for the node you want to break, and choose **Break Node**.

Figure 215: Break a Node

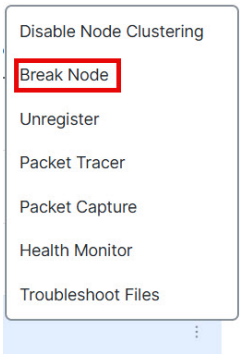

You can optionally break one or more nodes from the cluster More menu by choosing **Break Nodes**.

Step 2 You are prompted to confirm the break; click **Yes**.

Figure 216: Confirm Break

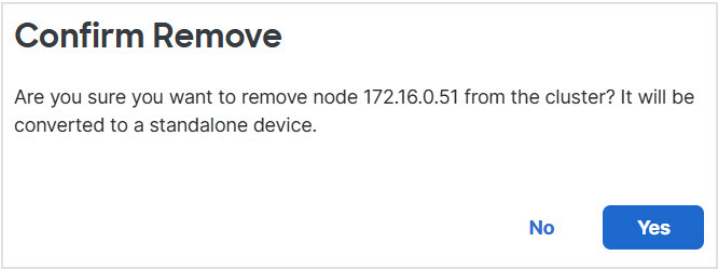

You can monitor the cluster node break by clicking the **Notifications** icon and choosing **Tasks**.

Break the Cluster

You can break the cluster and convert all nodes to standalone devices. The control node retains the interface and security policy configuration, while data nodes have their configuration erased.

Procedure

Step 1 Make sure all cluster nodes are being managed by the Firewall Management Center by reconciling nodes. See Reconcile Cluster Nodes, on page 497.

Step 2 Choose **Devices** > **Device Management**, click **More** (⋮) for the cluster, and choose **Break Cluster**.

Figure 217: Break Cluster

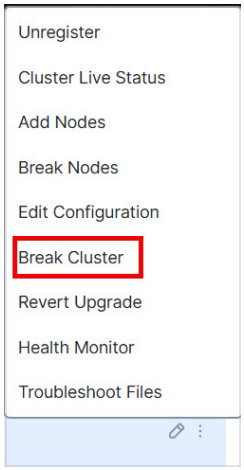

Step 3 You are prompted to break the cluster; click **Yes**.

Figure 218: Confirm Break

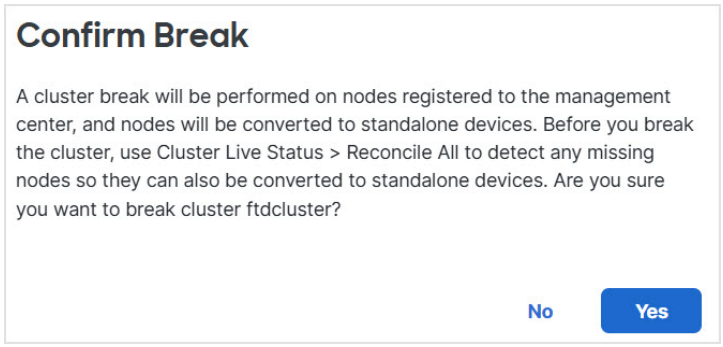

You can monitor the cluster break by clicking the **Notifications** icon and choosing **Tasks**.

Disable Clustering

You may want to deactivate a node in preparation for deleting the node, or temporarily for maintenance. This procedure is meant to temporarily deactivate a node; the node will still appear in the Firewall Management Center device list. When a node becomes inactive, all data interfaces are shut down.

Procedure

Step 1 For the unit you want to disable, choose **Devices** > **Device Management**, click **More** (⋮), and choose **Disable Node Clustering**.

Figure 219: Disable Clustering

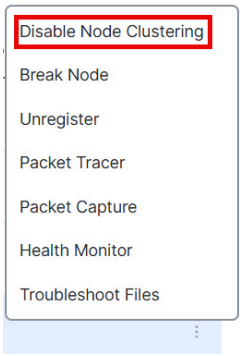

If you disable clustering on the control node, one of the data nodes will become the new control node. Note that for centralized features, if you force a control node change, then all connections are dropped, and you have to re-establish the connections on the new control node. You cannot disable clustering on the control node if it is the only node in the cluster.

Step 2 Confirm that you want to disable clustering on the node.

The node will show **(Disabled)** next to its name in the **Devices** > **Device Management** list.

Step 3 To reenable clustering, see Rejoin the Cluster, on page 494.

Rejoin the Cluster

If a node was removed from the cluster, for example for a failed interface or if you manually disabled clustering, you must manually rejoin the cluster. Make sure the failure is resolved before you try to rejoin the cluster. See Rejoining the Cluster, on page 517 for more information about why a node can be removed from a cluster.

Procedure

Step 1 For the unit you want to reactivate, choose **Devices** > **Device Management**, click **More** (⋮), and choose **Enable Node Clustering**.

Step 2 Confirm that you want to enable clustering on the unit.

Change the Control Node

⚠

Caution The best method to change the control node is to disable clustering on the control node, wait for a new control election, and then re-enable clustering. If you must specify the *exact* unit you want to become the control node, use the procedure in this section. Note that for centralized features, if you force a control node change using either method, then all connections are dropped, and you have to re-establish the connections on the new control node.

To change the control node, perform the following steps.

Procedure

Step 1 Open the **Cluster Status** dialog box by choosing **Devices** > **Device Management More** (⋮) **Cluster Live Status**.

Figure 220: Cluster Status

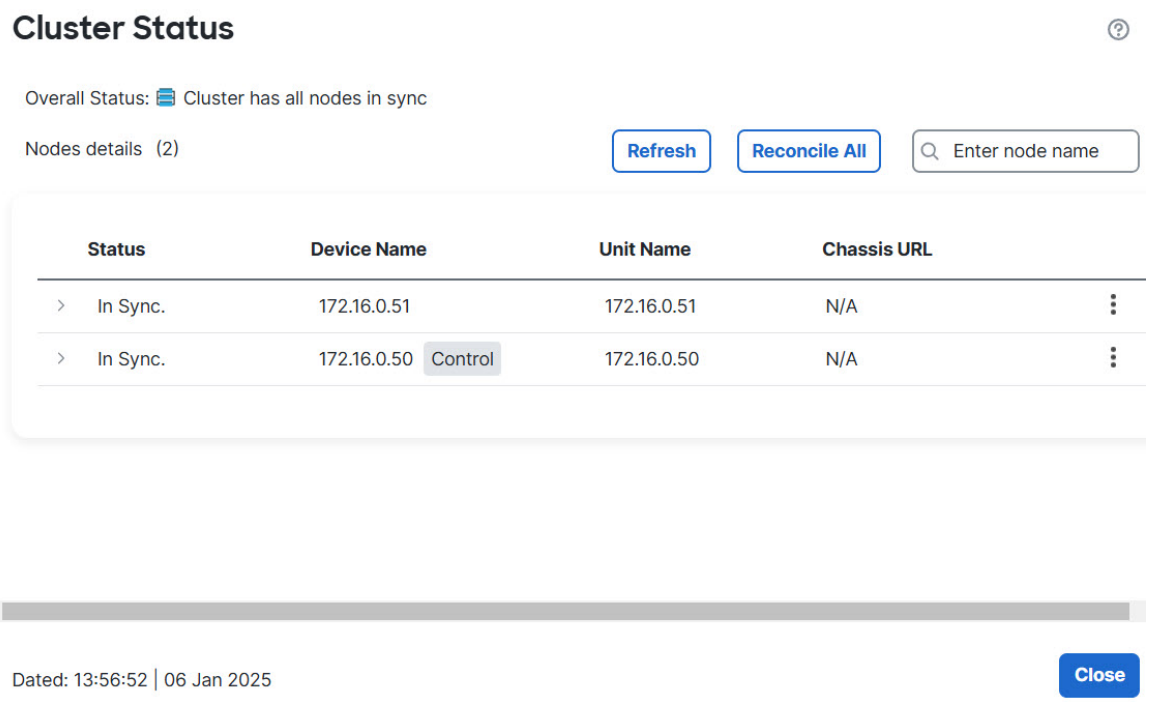

Step 2 For the unit you want to become the control unit, choose **More** (⋮) **Change Role to Control**.

Step 3 You are prompted to confirm the role change. Check the checkbox, and click **OK**.

Edit the Cluster Configuration

You can edit the cluster configuration. If you change the cluster key, cluster control link interface, or cluster control link network, the cluster will be broken and reformed automatically. Until the cluster is reformed, you may experience traffic disruption. If you change the cluster control link IP address for a node, node priority, or site ID, only the affected nodes are broken and readded to the cluster.

Procedure

Step 1 Choose **Devices** > **Device Management**, click **More** (⋮) for the cluster, and choose **Edit Configuration**.

Figure 221: Edit Configuration

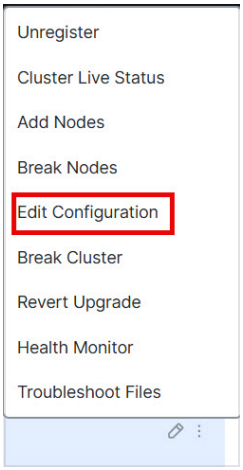

The **Manage Cluster Wizard** appears.

Step 2 Update the cluster configuration.

Figure 222: Manage Cluster Wizard

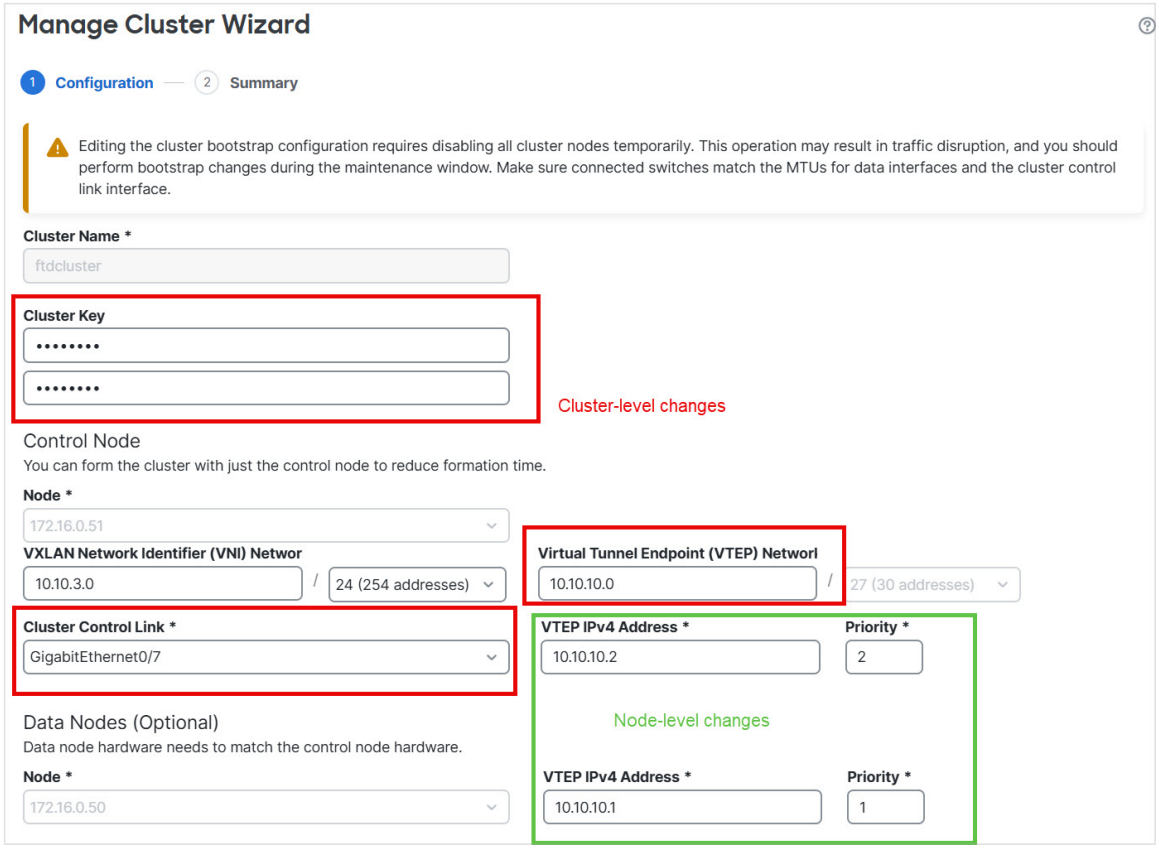

If the cluster control link is an EtherChannel, you can edit the interface membership and LACP configuration by clicking **Edit** (✎) next to the interface drop-down menu.

Step 3 Click **Continue**. Review the **Summary**, and then click **Save**

Reconcile Cluster Nodes

If a cluster node fails to register, you can reconcile the cluster membership from the device to the Firewall Management Center. For example, a data node might fail to register if the Firewall Management Center is occupied with certain processes, or if there is a network issue.

Procedure

Step 1 Choose **Devices** > **Device Management More** (⋮) for the cluster, and then choose **Cluster Live Status** to open the **Cluster Status** dialog box.

Figure 223: Cluster Live Status

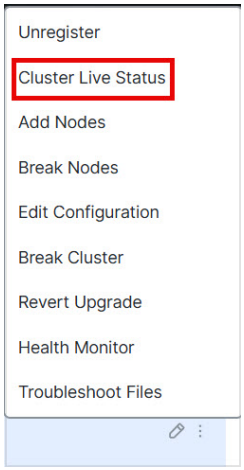

Step 2 Click **Reconcile All**.

Figure 224: Reconcile All

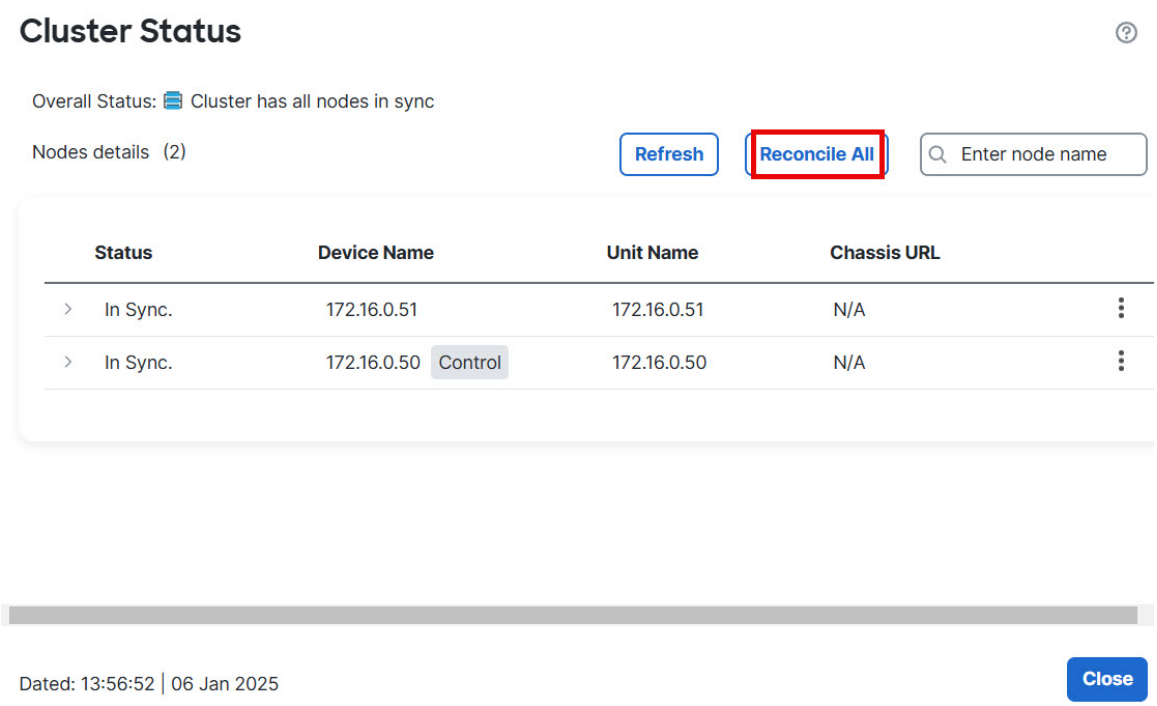

For more information about the cluster status, see Monitoring the Cluster, on page 499.

Unregister the Cluster or Nodes and Register to a New Firewall Management Center

You can unregister the cluster from the Firewall Management Center, which keeps the cluster intact. You might want to unregister the cluster if you want to add the cluster to a new Firewall Management Center.

You can also unregister a node from the Firewall Management Center without breaking the node from the cluster. Although the node is not visible in the Firewall Management Center, it is still part of the cluster, and it will continue to pass traffic and could even become the control node. You cannot unregister the current control node. You might want to unregister the node if it is no longer reachable from the Firewall Management Center, but you still want to keep it as part of the cluster while you troubleshoot management connectivity.

Unregistering a cluster:

- Severs all communication between the Firewall Management Center and the cluster.
- Removes the cluster from the **Device Management** page.
- Returns the cluster to local time management if the cluster's platform settings policy is configured to receive time from the Firewall Management Center using NTP.
- Leaves the configuration intact, so the cluster continues to process traffic.

  Policies, such as NAT and VPN, ACLs, and the interface configurations remain intact.

Registering the cluster again to the same or a different Firewall Management Center causes the configuration to be removed, so the cluster will stop processing traffic at that point; the cluster configuration remains intact so you can add the cluster as a whole. You can choose an access control policy at registration, but you will have to re-apply other policies after registration and then deploy the configuration before it will process traffic again.

Before you begin

This procedure requires CLI access to one of the nodes.

Procedure

Step 1 Choose **Devices** > **Device Management**, click **More** (⋮) for the cluster or node, and choose **Unregister**.

Figure 225: Unregister Cluster or Node

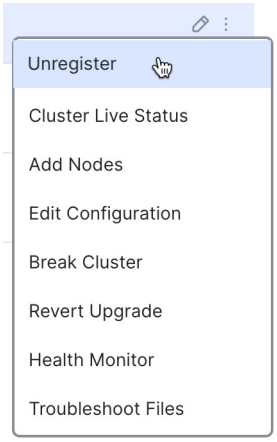

Step 2 You are prompted to unregister the cluster or node; click **Yes**.

Step 3 You can register the cluster to a new (or the same) Firewall Management Center by adding one of the cluster members as a new device.

You only need to add one of the cluster nodes as a device, and the rest of the cluster nodes will be discovered.

a) Connect to one cluster node's CLI, and identify the new Firewall Management Center using the **configure manager add** command. See Modify Firewall Threat Defense Management Interfaces at the CLI, on page 158.

b) Choose **Devices** > **Device Management**, and then click **Add** > **Device**.

Step 4 To re-add an unregistered node, see Reconcile Cluster Nodes, on page 497.

Monitoring the Cluster

You can monitor the cluster in the Firewall Management Center and at the Firewall Threat Defense CLI.

- **Cluster Status** dialog box, which is available from the **Devices** > **Device Management More** (⋮) icon or from the **Devices** > **Device Management** > **Cluster** page **General** area **Cluster Live Status** link.

Figure 226: Cluster Status

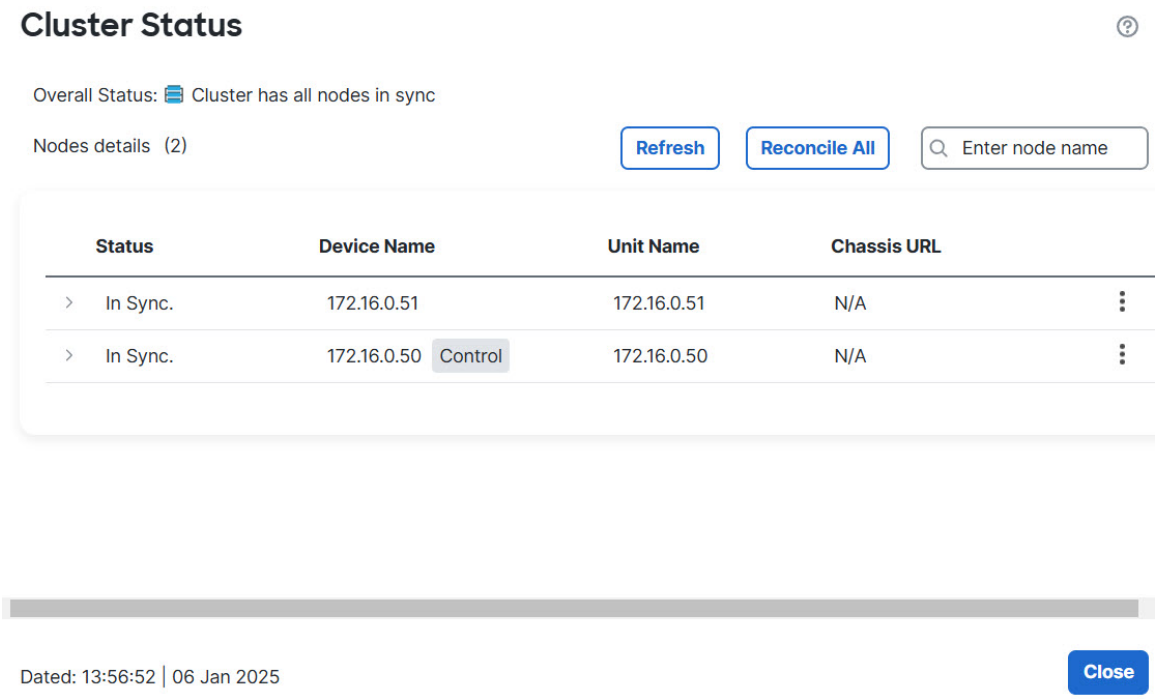

The Control node has a graphic indicator identifying its role.

Cluster member **Status** includes the following states:

- In Sync.—The node is registered with the Firewall Management Center.
- Pending Registration—The node is part of the cluster, but has not yet registered with the Firewall Management Center. If a node fails to register, you can retry registration by clicking **Reconcile All**.
- Clustering is disabled—The node is registered with the Firewall Management Center, but is an inactive member of the cluster. The clustering configuration remains intact if you intend to later re-enable it, or you can delete the node from the cluster.
- Joining cluster...—The node is joining the cluster on the chassis, but has not completed joining. After it joins, it will register with the Firewall Management Center.

For each node, you can view the **Summary** or the **History**.

Figure 227: Node Summary

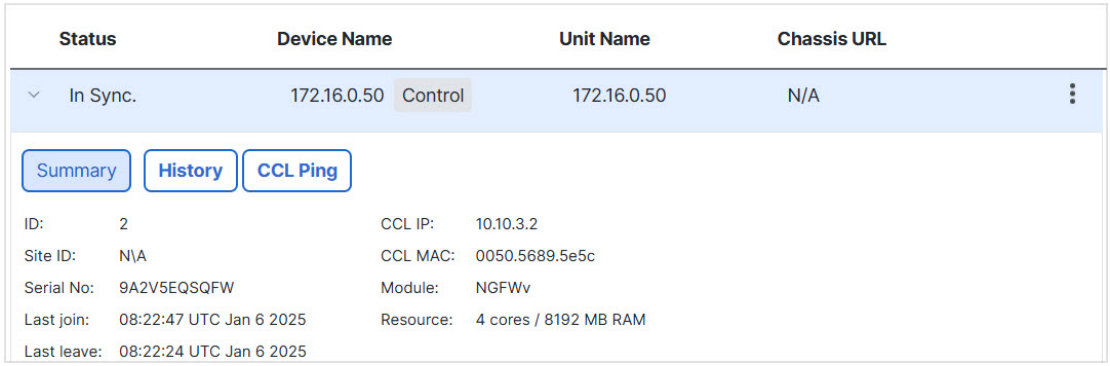

Figure 228: Node History

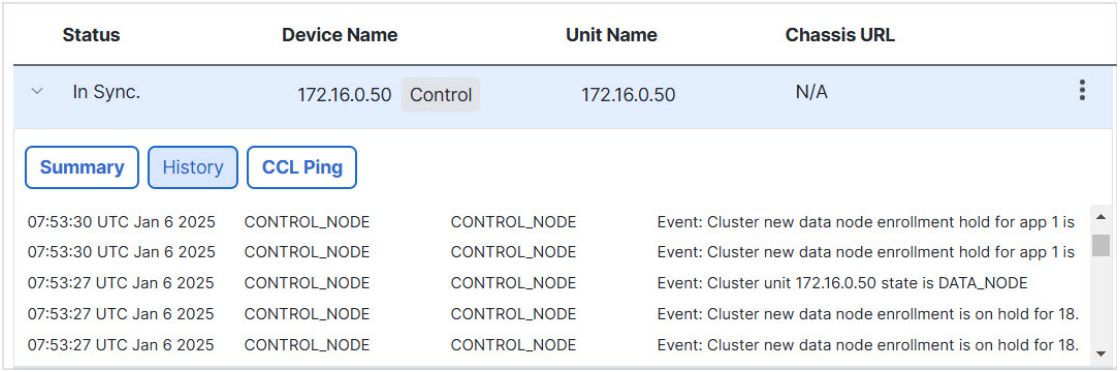

- **System** (⚙) > **Tasks** page.

  The **Tasks** page shows updates of the Cluster Registration task as each node registers.

- **Devices** > **Device Management***cluster_name*.

  When you expand the cluster on the devices listing page, you can see all member nodes, including the control node shown with its role next to the IP address. For nodes that are still registering, you can see the loading icon.

- **show cluster** {**access-list** [*acl_name*] | **conn** [**count**] | **cpu** [**usage**] | **history** | **interface-mode** | **memory** | **resource usage** | **service-policy** | **traffic** | **xlate count**}

  To view aggregated data for the entire cluster or other information, use the **show cluster** command.

- **show cluster info** [**auto-join** | **clients** | **conn-distribution** | **flow-mobility counters** | **goid** [*options*] | **health** | **incompatible-config** | **loadbalance** | **old-members** | **packet-distribution** | **trace** [*options*] | **transport** { **asp** | **cp**}]

  To view cluster information, use the **show cluster info** command.

Cluster Health Monitor Dashboard

Cluster Health Monitor

When a Firewall Threat Defense is the control node of a cluster, the Firewall Management Center collects various metrics periodically from the device metric data collector. The cluster health monitor is comprised of the following components:

- Overview dashboard—Displays information about the cluster topology, cluster statistics, and metric charts:
  - The topology section displays a cluster's live status, the health of individual threat defense, threat defense node type (control node or data node), and the status of the device. The status of the device could be *Disabled* (when the device leaves the cluster), *Added out of box* (in a public cloud cluster, the additional nodes that do not belong to the Firewall Management Center), or *Normal* (ideal state of the node).
  - The cluster statistics section displays current metrics of the cluster with respect to the CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.

Note The CPU and memory metrics display the individual average of the data plane and snort usage.

  - The metric charts, namely, CPU Usage, Memory Usage, Throughput, and Connections, diagrammatically display the statistics of the cluster over the specified time period.

- Load Distribution dashboard—Displays load distribution across the cluster nodes in two widgets:
  - The Distribution widget displays the average packet and connection distribution over the time range across the cluster nodes. This data depicts how the load is being distributed by the nodes. Using this widget, you can easily identify any abnormalities in the load distribution and rectify it.
  - The Node Statistics widget displays the node level metrics in table format. It displays metric data on CPU usage, memory usage, input rate, output rate, active connections, and NAT translations across the cluster nodes. This table view enables you to correlate data and easily identify any discrepancies.
- Member Performance dashboard—Displays current metrics of the cluster nodes. You can use the selector to filter the nodes and view the details of a specific node. The metric data include CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.
- CCL dashboard—Displays, graphically, the cluster control link data namely, the input, and output rate.
- Troubleshooting and Links — Provides convenient links to frequently used troubleshooting topics and procedures.
- Time range—An adjustable time window to constrain the information that appears in the various cluster metrics dashboards and widgets.
- Custom Dashboard—Displays data on both cluster-wide metrics and node-level metrics. However, node selection only applies for the threat defense metrics and not for the entire cluster to which the node belongs.

Viewing Cluster Health

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

The cluster health monitor provides a detailed view of the health status of a cluster and its nodes. This cluster health monitor provides health status and trends of the cluster in an array of dashboards.

Before you begin

- Ensure you have created a cluster from one or more devices in the Firewall Management Center.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Use the Monitoring navigation pane to access node-specific health monitors.

Step 2 In the device list, click **Expand**(>) and **Collapse** (∨) to expand and collapse the list of managed cluster devices.

Step 3 To view the cluster health statistics, click on the cluster name. The cluster monitor reports health and performance metrics in several predefined dashboards by default. The metrics dashboards include:

- Overview — Highlights key metrics from the other predefined dashboards, including its nodes, CPU, memory, input and output rates, connection statistics, and NAT translation information.
- Load Distribution — Traffic and packet distribution across the cluster nodes.
- Member Performance — Node-level statistics on CPU usage, memory usage, input throughput, output throughput, active connection, and NAT translation.
- CCL — Interface status and aggregate traffic statistics.

You can navigate through the various metrics dashboards by clicking on the labels. For a comprehensive list of the supported cluster metrics, see Cisco Secure Firewall Threat Defense Health Metrics.

Step 4 You can configure the time range from the drop-down in the upper-right corner. The time range can reflect a period as short as the last hour (the default) or as long as two weeks. Select **Custom** from the drop-down to configure a custom start and end date.

Click the refresh icon to set auto refresh to 5 minutes or to toggle off auto refresh.

Step 5 Click on deployment icon for a deployment overlay on the trend graph, with respect to the selected time range.

The deployment icon indicates the number of deployments during the selected time-range. A vertical band indicates the deployment start and end time. For multiple deployments, multiple bands/lines appear. Click on the icon on top of the dotted line to view the deployment details.

Step 6 (For node-specific health monitor) View the **Health Alerts** for the node in the alert notification at the top of page, directly to the right of the device name.

Hover your pointer over the **Health Alerts** to view the health summary of the node. The popup window shows a truncated summary of the top five health alerts. Click on the popup to open a detailed view of the health alert summary.

Step 7 (For node-specific health monitor) The device monitor reports health and performance metrics in several predefined dashboards by default. The metrics dashboards include:

- Overview — Highlights key metrics from the other predefined dashboards, including CPU, memory, interfaces, connection statistics; plus disk usage and critical process information.
- CPU — CPU utilization, including the CPU usage by process and by physical cores.
- Memory — Device memory utilization, including data plane and Snort memory usage.
- Interfaces — Interface status and aggregate traffic statistics.
- Connections — Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
- Snort — Statistics that are related to the Snort process.
- ASP drops — Statistics related to the dropped packets against various reasons.

You can navigate through the various metrics dashboards by clicking on the labels. See Cisco Secure Firewall Threat Defense Health Metrics for a comprehensive list of the supported device metrics.

Step 8 Click the plus sign **Add New Dashboard**(+) in the upper right corner of the health monitor to create a custom dashboard by building your own variable set from the available metric groups.

For cluster-wide dashboard, choose Cluster metric group, and then choose the metric.

Cluster Metrics

The cluster health monitor tracks statistics that are related to a cluster and its nodes, and aggregate of load distribution, performance, and CCL traffic statistics.

Table 32: Cluster Metrics

Metric	Description	Format
CPU	Average of CPU metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Memory	Average of memory metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Data Throughput	Incoming and outgoing data traffic statistics for a cluster.	bytes
CCL Throughput	Incoming and outgoing CCL traffic statistics for a cluster.	bytes
Connections	Count of active connections in a cluster.	number
NAT Translations	Count of NAT translations for a cluster.	number
Distribution	Connection distribution count in the cluster for every second.	number

Metric	Description	Format
Packets	Packet distribution count in the cluster for every second.	number

Troubleshooting the Cluster

You can use the **CCL Ping** tool to make sure the cluster control link is operating correctly. You can also use the following tools that are available for devices and clusters:

- Troubleshooting files—If a node fails to join the cluster, a troubleshooting file is automatically generated. You can also generate and download troubleshooting files from the **Devices** > **Device Management** > **Cluster** > **General** area. See Generate Troubleshooting Files, on page 129.

  You can also generate files from the **Device Management** page by clicking **More** (⋮) and choosing **Troubleshoot Files**.

- CLI output—From the **Devices** > **Device Management** > **Cluster** > **General** area, you can view a set of pre-defined CLI outputs that can help you troubleshoot the cluster. The following commands are automatically run for the cluster:

  - **show running-config cluster**
  - **show cluster info**
  - **show cluster info health**
  - **show cluster info transport cp**
  - **show version**
  - **show asp drop**
  - **show counters**
  - **show arp**
  - **show int ip brief**
  - **show blocks**
  - **show cpu detailed**
  - **show interface** *ccl_interface*
  - **ping** *ccl_ip* **size** *ccl_mtu* **repeat 2**

  You can also enter any **show** command in the Command field. See View CLI Output, on page 132 for more information.

Perform a Ping on the Cluster Control Link

When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the ping fails, a notification is generated so you can fix

the MTU mismatch on connecting switches and try again. This tool lets you manually ping all nodes that have already joined the cluster in case you are having cluster control link connectivity problems.

Procedure

Step 1 Choose **Devices** > **Device Management**, click the **More** (⋮) icon next to the cluster, and choose **Cluster Live Status**.

Figure 229: Cluster Status

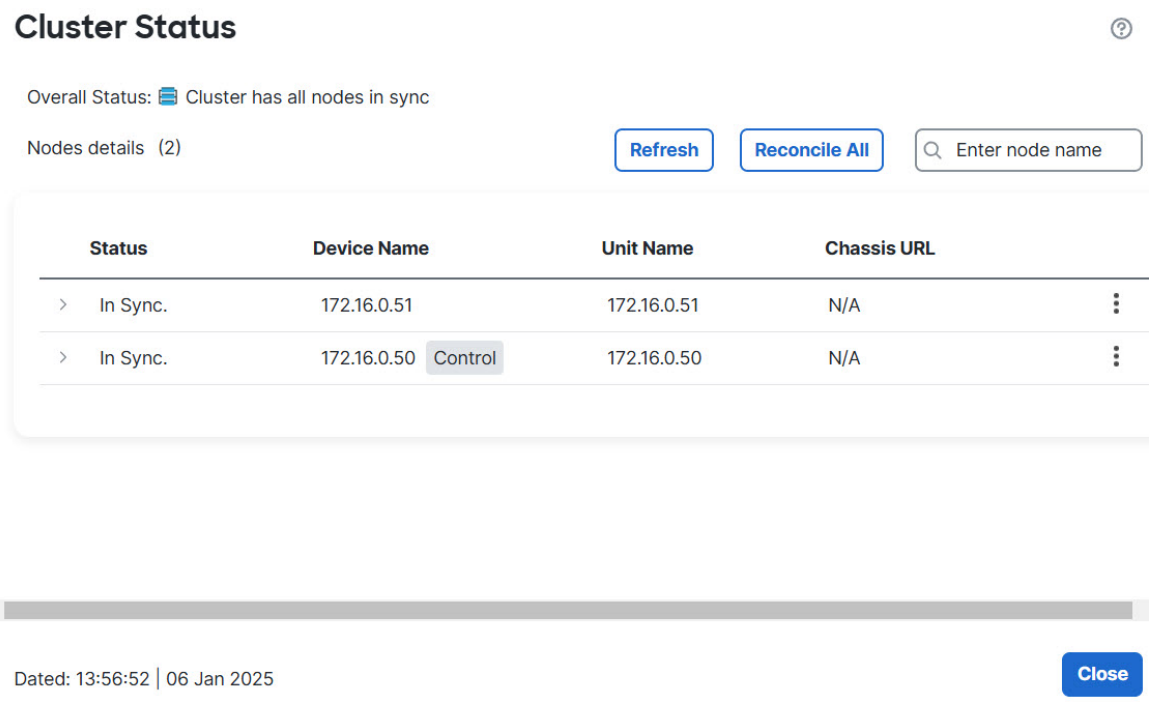

Step 2 Expand one of the nodes, and click **CCL Ping**.

Figure 230: CCL Ping

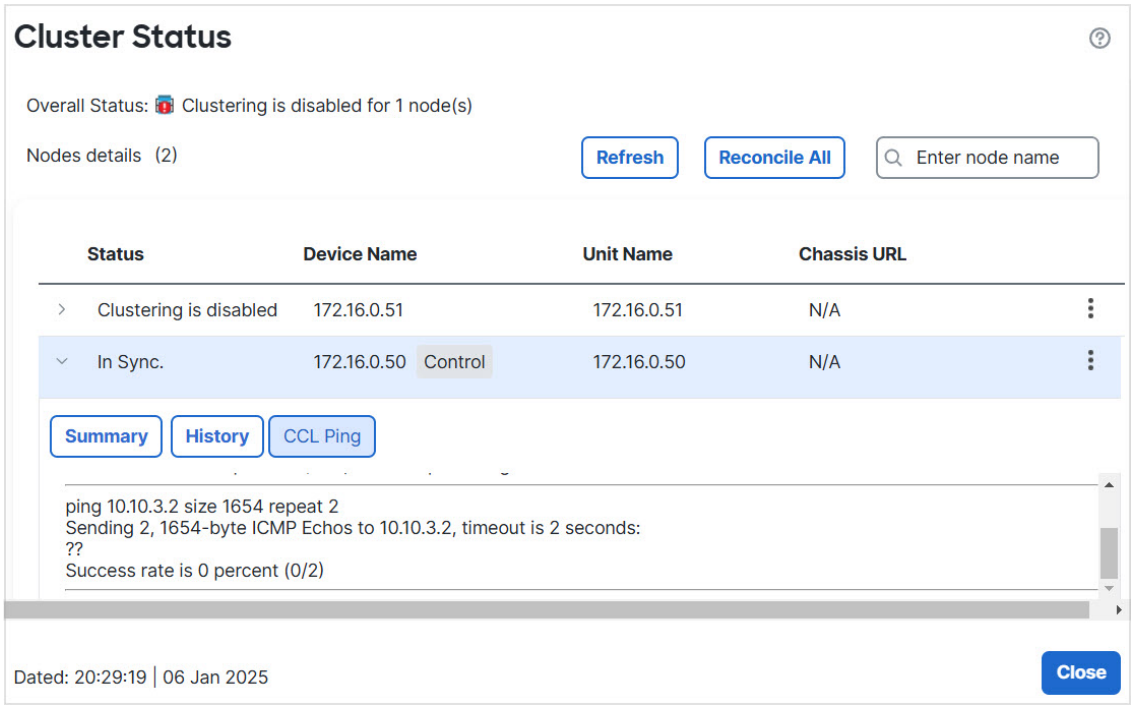

The node sends a ping on the cluster control link to every other node using a packet size that matches the maximum MTU.

Examples for Clustering

These examples include examples for typical deployments.

Firewall on a Stick

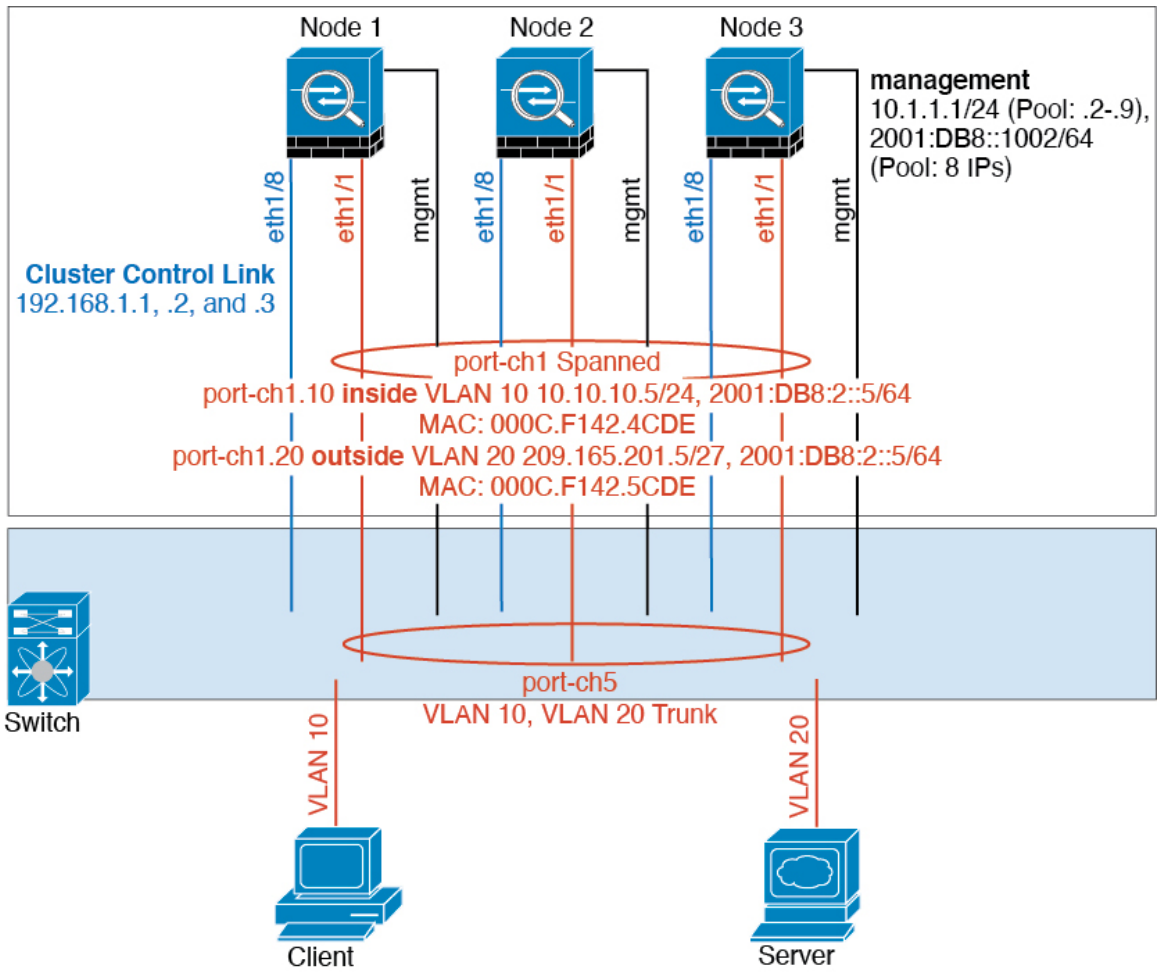

Data traffic from different security domains are associated with different VLANs, for example, VLAN 10 for the inside network and VLAN 20 for the outside network. Each has a single physical port connected to the external switch or router. Trunking is enabled so that all packets on the physical link are 802.1q encapsulated. The is the firewall between VLAN 10 and VLAN 20.

When using Spanned EtherChannels, all data links are grouped into one EtherChannel on the switch side. If the becomes unavailable, the switch will rebalance traffic between the remaining units.

Traffic Segregation

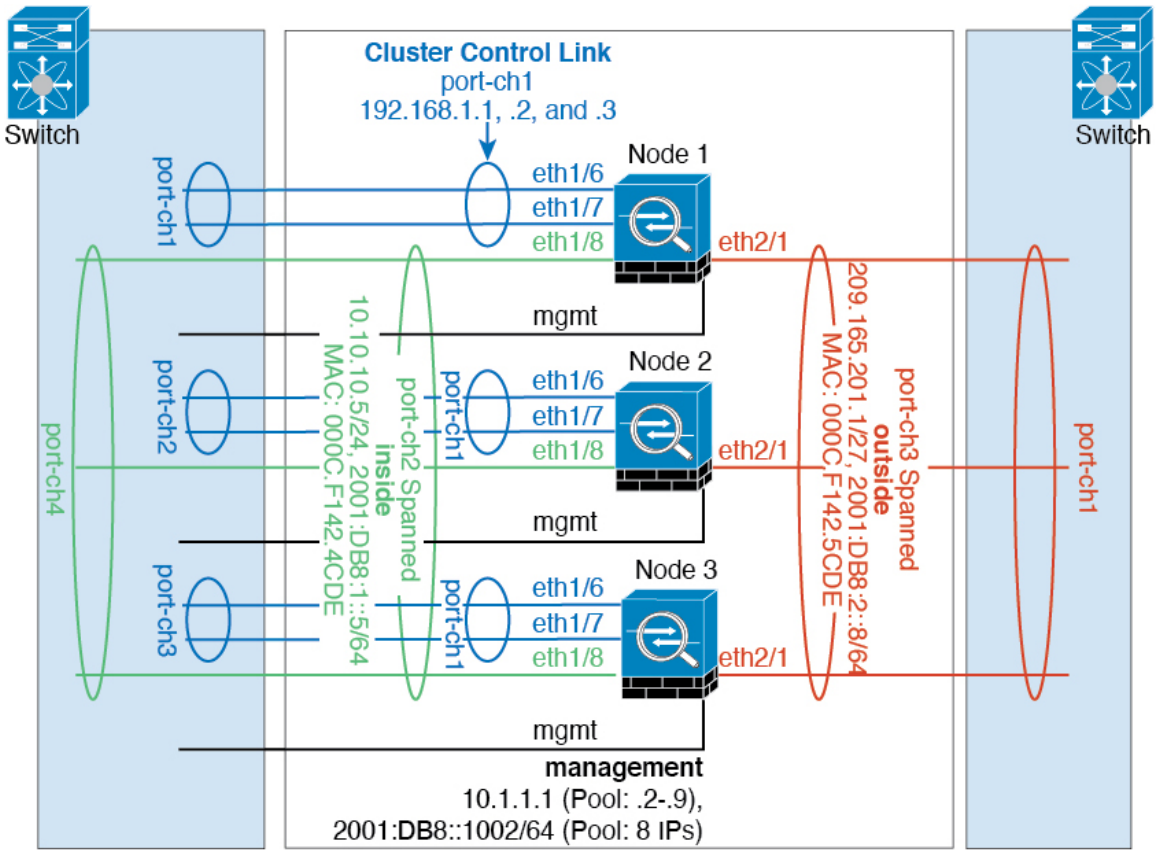

You may prefer physical separation of traffic between the inside and outside network.

As shown in the diagram above, there is one Spanned EtherChannel on the left side that connects to the inside switch, and the other on the right side to outside switch. You can also create VLAN subinterfaces on each EtherChannel if desired.

Reference for Clustering

This section includes more information about how clustering operates.

Firewall Threat Defense Features and Clustering

Some Firewall Threat Defense features are not supported with clustering, and some are only supported on the control unit. Other features might have caveats for proper usage.

Unsupported Features with Clustering

These features cannot be configured with clustering enabled, and the commands will be rejected.

Note To view FlexConfig features that are also not supported with clustering, for example WCCP inspection, see the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the Firewall Management Center GUI. See FlexConfig Policies, on page 2667.

- Remote access VPN (SSL VPN and IPsec VPN)
- DHCP client, server, and proxy. DHCP relay is supported.
- Virtual Tunnel Interfaces (VTIs)
- High Availability
- Integrated Routing and Bridging
- Firewall Management Center UCAPL/CC mode
- DHCP client, server, and proxy. DHCP relay is supported.

Centralized Features for Clustering

The following features are only supported on the control node, and are not scaled for the cluster.

Note Traffic for centralized features is forwarded from member nodes to the control node over the cluster control link.

If you use the rebalancing feature, traffic for centralized features may be rebalanced to non-control nodes before the traffic is classified as a centralized feature; if this occurs, the traffic is then sent back to the control node.

For centralized features, if the control node fails, all connections are dropped, and you have to re-establish the connections on the new control node.

Note To view FlexConfig features that are also centralized with clustering, for example RADIUS inspection, see the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the Firewall Management Center GUI. See FlexConfig Policies, on page 2667.

- The following application inspections:
  - DCERPC
  - ESMTP
  - NetBIOS
  - PPTP
  - RSH
  - SQLNET

- SUNRPC
- TFTP
- XDMCP

- Static route monitoring
- Site-to-site VPN
- IGMP multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- PIM multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- Dynamic routing (Spanned EtherChannel mode only)

Connection Settings and Clustering

Connection limits are enforced cluster-wide. Each node has an estimate of the cluster-wide counter values based on broadcast messages. Due to efficiency considerations, the configured connection limit across the cluster might not be enforced exactly at the limit number. Each node may overestimate or underestimate the cluster-wide counter value at any given time. However, the information will get updated over time in a load-balanced cluster.

FTP and Clustering

- If FTP data channel and control channel flows are owned by different cluster members, then the data channel owner will periodically send idle timeout updates to the control channel owner and update the idle timeout value. However, if the control flow owner is reloaded, and the control flow is re-hosted, the parent/child flow relationship will not longer be maintained; the control flow idle timeout will not be updated.

Multicast Routing in Individual Interface Mode

In Individual interface mode, units do not act independently with multicast. All data and routing packets are processed and forwarded by the control unit, thus avoiding packet replication.

Multicast Routing in Individual Interface Mode

In Individual interface mode, units do not act independently with multicast. All data and routing packets are processed and forwarded by the control unit, thus avoiding packet replication.

NAT and Clustering

NAT can affect the overall throughput of the cluster. Inbound and outbound NAT packets can be sent to different Firewall Threat Defenses in the cluster, because the load balancing algorithm relies on IP addresses and ports, and NAT causes inbound and outbound packets to have different IP addresses and/or ports. When a packet arrives at the Firewall Threat Defense that is not the NAT owner, it is forwarded over the cluster control link to the owner, causing large amounts of traffic on the cluster control link. Note that the receiving node does not create a forwarding flow to the owner, because the NAT owner may not end up creating a connection for the packet depending on the results of security and policy checks.

If you still want to use NAT in clustering, then consider the following guidelines:

- No Proxy ARP—For Individual interfaces, a proxy ARP reply is never sent for mapped addresses. This prevents the adjacent router from maintaining a peer relationship with a node that may no longer be in the cluster. The upstream router needs a static route or PBR with Object Tracking for the mapped addresses that points to the Main cluster IP address. This is not an issue for a Spanned EtherChannel, because there is only one IP address associated with the cluster interface.
- No interface PAT on an Individual interface—Interface PAT is not supported for Individual interfaces.
- PAT with Port Block Allocation—See the following guidelines for this feature:
  - Maximum-per-host limit is not a cluster-wide limit, and is enforced on each node individually. Thus, in a 3-node cluster with the maximum-per-host limit configured as 1, if the traffic from a host is load-balanced across all 3 nodes, then it can get allocated 3 blocks with 1 in each node.
  - Port blocks created on the backup node from the backup pools are not accounted for when enforcing the maximum-per-host limit.
  - On-the-fly PAT rule modifications, where the PAT pool is modified with a completely new range of IP addresses, will result in xlate backup creation failures for the xlate backup requests that were still in transit while the new pool became effective. This behavior is not specific to the port block allocation feature, and is a transient PAT pool issue seen only in cluster deployments where the pool is distributed and traffic is load-balanced across the cluster nodes.
  - When operating in a cluster, you cannot simply change the block allocation size. The new size is effective only after you reload each device in the cluster. To avoid having to reload each device, we recommend that you delete all block allocation rules and clear all xlates related to those rules. You can then change the block size and recreate the block allocation rules.
- NAT pool address distribution for dynamic PAT—When you configure a PAT pool, the cluster divides each IP address in the pool into port blocks. By default, each block is 512 ports, but if you configure port block allocation rules, your block setting is used instead. These blocks are distributed evenly among the nodes in the cluster, so that each node has one or more blocks for each IP address in the PAT pool. Thus, you could have as few as one IP address in a PAT pool for a cluster, if that is sufficient for the number of PAT'ed connections you expect. Port blocks cover the 1024-65535 port range, unless you configure the option to include the reserved ports, 1-1023, on the PAT pool NAT rule.
- Reusing a PAT pool in multiple rules—To use the same PAT pool in multiple rules, you must be careful about the interface selection in the rules. You must either use specific interfaces in all rules, or "any" in all rules. You cannot mix specific interfaces and "any" across the rules, or the system might not be able to match return traffic to the right node in the cluster. Using unique PAT pools per rule is the most reliable option.
- No round-robin—Round-robin for a PAT pool is not supported with clustering.
- No extended PAT—Extended PAT is not supported with clustering.
- Dynamic NAT xlates managed by the control node—The control node maintains and replicates the xlate table to data nodes. When a data node receives a connection that requires dynamic NAT, and the xlate is not in the table, it requests the xlate from the control node. The data node owns the connection.
- Stale xlates—The xlate idle time on the connection owner does not get updated. Thus, the idle time might exceed the idle timeout. An idle timer value higher than the configured timeout with a refcnt of 0 is an indication of a stale xlate.
- No static PAT for the following inspections—

- FTP
- RSH
- SQLNET
- TFTP
- XDMCP
- SIP

- If you have an extremely large number of NAT rules, over ten thousand, you should enable the transactional commit model using the **asp rule-engine transactional-commit nat** command in the device CLI. Otherwise, the node might not be able to join the cluster.

Dynamic Routing

The routing process only runs on the control node, and routes are learned through the control node and replicated to data nodes. If a routing packet arrives at a data node, it is redirected to the control node.

Figure 231: Dynamic Routing in Spanned EtherChannel Mode

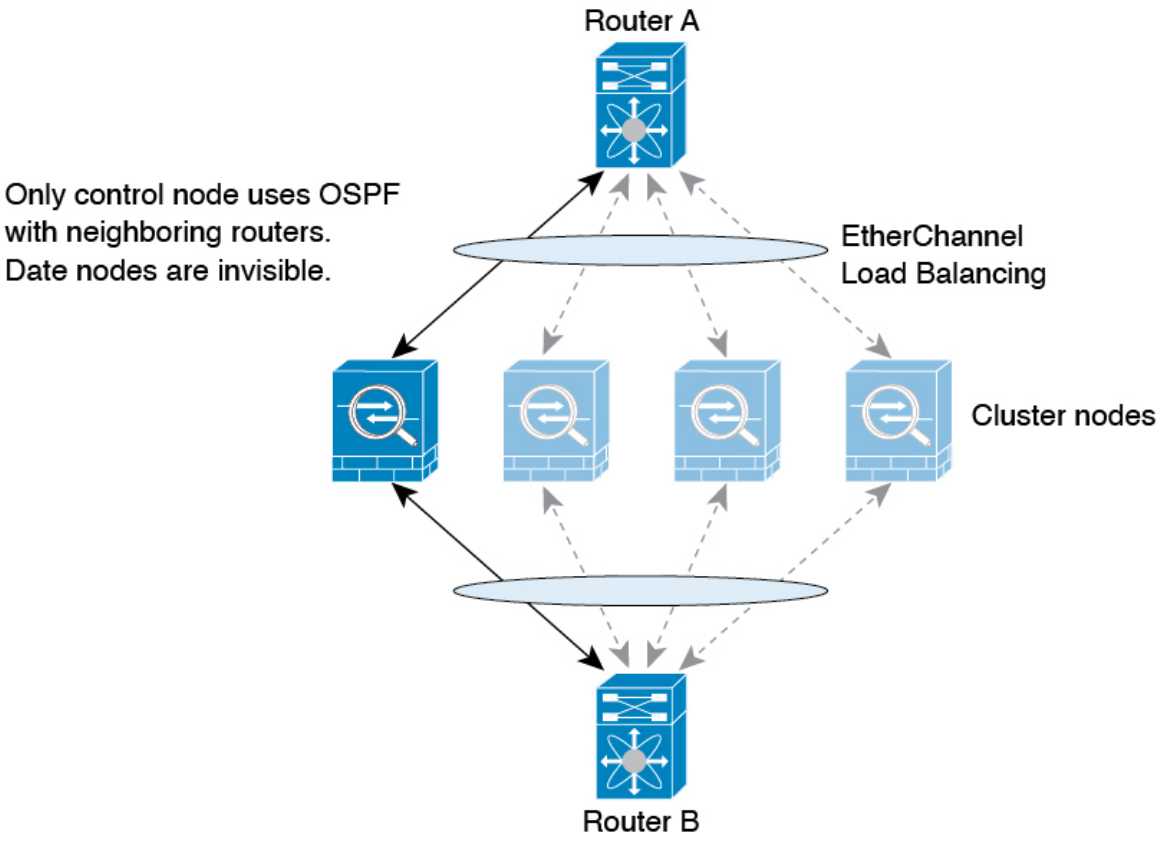

After the data node learn the routes from the control node, each node makes forwarding decisions independently.

The OSPF LSA database is not synchronized from the control node to data nodes. If there is a control node switchover, the neighboring router will detect a restart; the switchover is not transparent. The OSPF process

picks an IP address as its router ID. Although not required, you can assign a static router ID to ensure a consistent router ID is used across the cluster. See the OSPF Non-Stop Forwarding feature to address the interruption.

Dynamic Routing in Individual Interface Mode

In Individual interface mode, each node runs the routing protocol as a standalone router, and routes are learned by each node independently.

Figure 232: Dynamic Routing in Individual Interface Mode

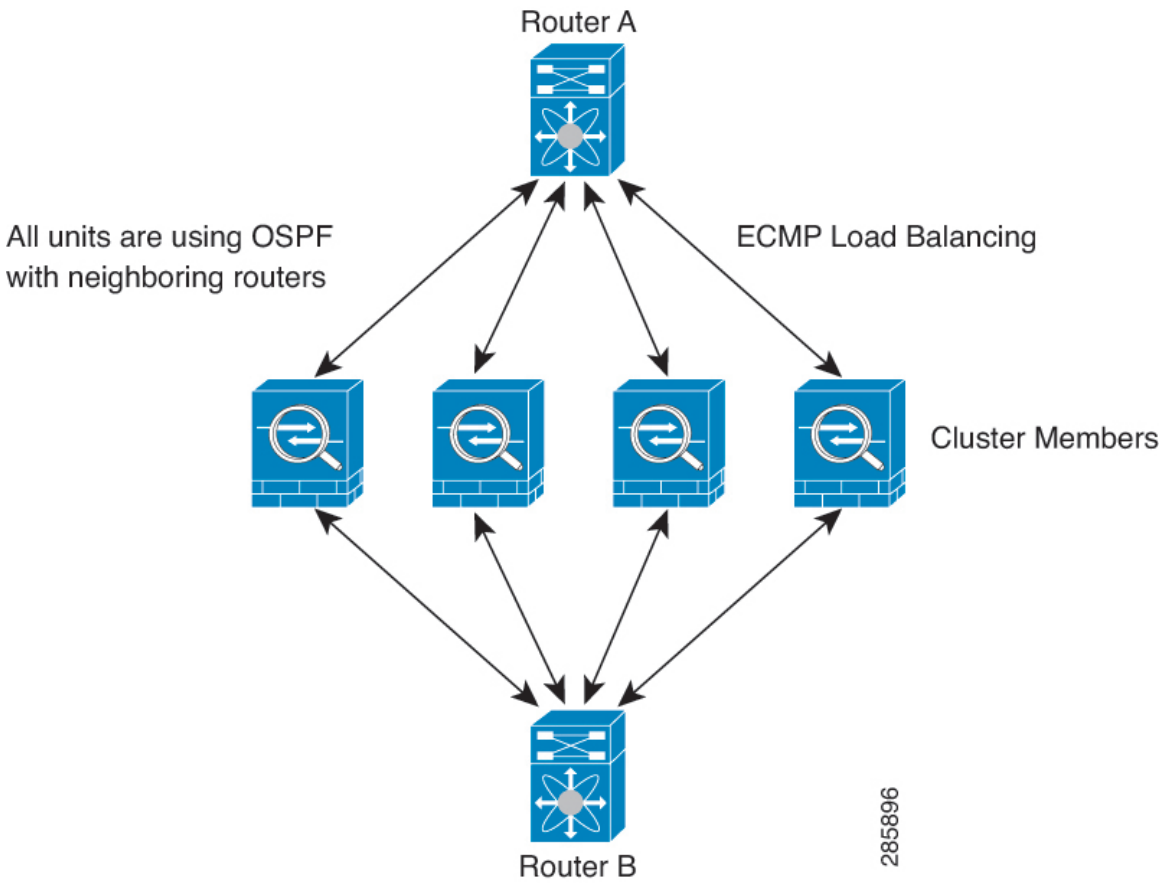

In the above diagram, Router A learns that there are 4 equal-cost paths to Router B, each through a node. ECMP is used to load balance traffic between the 4 paths. Each node picks a different router ID when talking to external routers.

You must configure a cluster pool for the router ID so that each node has a separate router ID.

EIGRP does not form neighbor relationships with cluster peers in individual interface mode.

Note If the cluster has multiple adjacencies to the same router for redundancy purposes, asymmetric routing can lead to unacceptable traffic loss. To avoid asymmetric routing, group all of these node interfaces into the same traffic zone. See Create an ECMP Zone, on page 1225.

SIP Inspection and Clustering

A control flow can be created on any node (due to load balancing); its child data flows must reside on the same node.

SNMP and Clustering

You should always use the Local address, and not the Main cluster IP address for SNMP polling. If the SNMP agent polls the Main cluster IP address, if a new control node is elected, the poll to the new control node will fail.

When using SNMPv3 with clustering, if you add a new cluster node after the initial cluster formation, then SNMPv3 users are not replicated to the new node. You must remove the users, and re-add them, and then redeploy your configuration to force the users to replicate to the new node.

Syslog and Clustering

- Each node in the cluster generates its own syslog messages. You can configure logging so that each node uses either the same or a different device ID in the syslog message header field. For example, the hostname configuration is replicated and shared by all nodes in the cluster. If you configure logging to use the hostname as the device ID, syslog messages generated by all nodes look as if they come from a single node. If you configure logging to use the local-node name that is assigned in the cluster bootstrap configuration as the device ID, syslog messages look as if they come from different nodes.

Cisco TrustSec and Clustering

Only the control node learns security group tag (SGT) information. The control node then populates the SGT to data nodes, and data nodes can make a match decision for SGT based on the security policy.

VPN and Clustering

Site-to-site VPN is a centralized feature; only the control node supports VPN connections.

Note Remote access VPN is not supported with clustering.

VPN functionality is limited to the control node and does not take advantage of the cluster high availability capabilities. If the control node fails, all existing VPN connections are lost, and VPN users will see a disruption in service. When a new control node is elected, you must reestablish the VPN connections.

When you connect a VPN tunnel to a Spanned EtherChannel address, connections are automatically forwarded to the control node. For connections to an Individual interface when using PBR or ECMP, you must always connect to the Main cluster IP address, not a Local address.

VPN-related keys and certificates are replicated to all nodes.

Performance Scaling Factor

When you combine multiple units into a cluster, you can expect the total cluster performance to be approximately 80% of the maximum combined throughput.

For example, if your model can handle approximately 10 Gbps of traffic when running alone, then for a cluster of 8 units, the maximum combined throughput will be approximately 80% of 80 Gbps (8 units x 10 Gbps): 64 Gbps.

Control Node Election

Nodes of the cluster communicate over the cluster control link to elect a control node as follows:

1. When you enable clustering for a node (or when it first starts up with clustering already enabled), it broadcasts an election request every 3 seconds.
2. Any other nodes with a higher priority respond to the election request; the priority is set between 1 and 100, where 1 is the highest priority.
3. If after 45 seconds, a node does not receive a response from another node with a higher priority, then it becomes the control node.

Note If multiple nodes tie for the highest priority, the cluster node name and then the serial number is used to determine the control node.

4. If a node later joins the cluster with a higher priority, it does not automatically become the control node; the existing control node always remains as the control node unless it stops responding, at which point a new control node is elected.
5. In a "split brain" scenario when there are temporarily multiple control nodes, then the node with highest priority retains the role while the other nodes return to data node roles.

Note You can manually force a node to become the control node. For centralized features, if you force a control node change, then all connections are dropped, and you have to re-establish the connections on the new control node.

High Availability Within the Cluster

Clustering provides high availability by monitoring node and interface health and by replicating connection states between nodes.

Node Health Monitoring

Each node periodically sends a broadcast heartbeat packet over the cluster control link. If the control node does not receive any heartbeat packets or other packets from a data node within the configurable timeout period, then the control node removes the data node from the cluster. If the data nodes do not receive packets from the control node, then a new control node is elected from the remaining nodes.

If nodes cannot reach each other over the cluster control link because of a network failure and not because a node has actually failed, then the cluster may go into a "split brain" scenario where isolated data nodes will elect their own control nodes. For example, if a router fails between two cluster locations, then the original control node at location 1 will remove the location 2 data nodes from the cluster. Meanwhile, the nodes at location 2 will elect their own control node and form their own cluster. Note that asymmetric traffic may fail

in this scenario. After the cluster control link is restored, then the control node that has the higher priority will keep the control node's role.

See Control Node Election, on page 516 for more information.

Interface Monitoring

Each node monitors the link status of all named hardware interfaces in use, and reports status changes to the control node.

- Spanned EtherChannel—Uses cluster Link Aggregation Control Protocol (cLACP). Each node monitors the link status and the cLACP protocol messages to determine if the port is still active in the EtherChannel. The status is reported to the control node.
- Individual interfaces (Routed mode only)—Each node self-monitors its interfaces and reports interface status to the control node.

When you enable health monitoring, all physical interfaces (including the main EtherChannel) are monitored by default; you can optionally disable monitoring per interface. Only named interfaces can be monitored. For example, the named EtherChannel must fail to be considered failed, which means all member ports of an EtherChannel must fail to trigger cluster removal.

A node is removed from the cluster if its monitored interfaces fail. The amount of time before the Firewall Threat Defense removes a member from the cluster depends on the type of interface and whether the node is an established member or is joining the cluster. For EtherChannels (spanned or not): If the interface is down on an established member, then the Firewall Threat Defense removes the member after 9 seconds. The Firewall Threat Defense does not monitor interfaces for the first 90 seconds that a node joins the cluster. Interface status changes during this time will not cause the Firewall Threat Defense to be removed from the cluster. For non-EtherChannels, the node is removed after 500 ms, regardless of the member state.

Status After Failure

When a node in the cluster fails, the connections hosted by that node are seamlessly transferred to other nodes; state information for traffic flows is shared over the control node's cluster control link.

If the control node fails, then another member of the cluster with the highest priority (lowest number) becomes the control node.

The Firewall Threat Defense automatically tries to rejoin the cluster, depending on the failure event.

Note When the Firewall Threat Defense becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the Management interface can send and receive traffic.

Rejoining the Cluster

After a cluster member is removed from the cluster, how it can rejoin the cluster depends on why it was removed:

- Failed cluster control link when initially joining—After you resolve the problem with the cluster control link, you must manually rejoin the cluster by re-enabling clustering.
- Failed cluster control link after joining the cluster—The Firewall Threat Defense automatically tries to rejoin every 5 minutes, indefinitely.

- Failed data interface—The Firewall Threat Defense automatically tries to rejoin at 5 minutes, then at 10 minutes, and finally at 20 minutes. If the join is not successful after 20 minutes, then the Firewall Threat Defense application disables clustering. After you resolve the problem with the data interface, you have to manually enable clustering.
- Failed node—If the node was removed from the cluster because of a node health check failure, then rejoining the cluster depends on the source of the failure. For example, a temporary power failure means the node will rejoin the cluster when it starts up again as long as the cluster control link is up. The Firewall Threat Defense application attempts to rejoin the cluster every 5 seconds.
- Internal error—Internal failures include: application sync timeout; inconsistent application statuses; and so on.
- Failed configuration deployment—If you deploy a new configuration from Firewall Management Center, and the deployment fails on some cluster members but succeeds on others, then the nodes that failed are removed from the cluster. You must manually rejoin the cluster by re-enabling clustering. If the deployment fails on the control node, then the deployment is rolled back, and no members are removed. If the deployment fails on all data nodes, then the deployment is rolled back, and no members are removed.

Data Path Connection State Replication

Every connection has one owner and at least one backup owner in the cluster. The backup owner does not take over the connection in the event of a failure; instead, it stores TCP/UDP state information, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner is usually also the director.

Some traffic requires state information above the TCP or UDP layer. See the following table for clustering support or lack of support for this kind of traffic.

Table 33: Features Replicated Across the Cluster

Traffic	State Support	Notes
Up time	Yes	Keeps track of the system up time.
ARP Table	Yes	—
MAC address table	Yes	—
User Identity	Yes	—
IPv6 Neighbor database	Yes	—
Dynamic routing	Yes	—
SNMP Engine ID	**No**	—

How the Cluster Manages Connections

Connections can be load-balanced to multiple nodes of the cluster. Connection roles determine how connections are handled in both normal operation and in a high availability situation.

Connection Roles

See the following roles defined for each connection:

- Owner—Usually, the node that initially receives the connection. The owner maintains the TCP state and processes packets. A connection has only one owner. If the original owner fails, then when new nodes receive packets from the connection, the director chooses a new owner from those nodes.
- Backup owner—The node that stores TCP/UDP state information received from the owner, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner does not take over the connection in the event of a failure. If the owner becomes unavailable, then the first node to receive packets from the connection (based on load balancing) contacts the backup owner for the relevant state information so it can become the new owner.

  As long as the director (see below) is not the same node as the owner, then the director is also the backup owner. If the owner chooses itself as the director, then a separate backup owner is chosen.

  For clustering on the Firepower 9300, which can include up to 3 cluster nodes in one chassis, if the backup owner is on the same chassis as the owner, then an additional backup owner will be chosen from another chassis to protect flows from a chassis failure.
- Director—The node that handles owner lookup requests from forwarders. When the owner receives a new connection, it chooses a director based on a hash of the source/destination IP address and ports (see below for ICMP hash details), and sends a message to the director to register the new connection. If packets arrive at any node other than the owner, the node queries the director about which node is the owner so it can forward the packets. A connection has only one director. If a director fails, the owner chooses a new director.

  As long as the director is not the same node as the owner, then the director is also the backup owner (see above). If the owner chooses itself as the director, then a separate backup owner is chosen.

  ICMP/ICMPv6 hash details:

  - For Echo packets, the source port is the ICMP identifier, and the destination port is 0.
  - For Reply packets, the source port is 0, and the destination port is the ICMP identifier.
  - For other packets, both source and destination ports are 0.
- Forwarder—A node that forwards packets to the owner. If a forwarder receives a packet for a connection it does not own, it queries the director for the owner, and then establishes a flow to the owner for any other packets it receives for this connection. The director can also be a forwarder. Note that if a forwarder receives the SYN-ACK packet, it can derive the owner directly from a SYN cookie in the packet, so it does not need to query the director. (If you disable TCP sequence randomization, the SYN cookie is not used; a query to the director is required.) For short-lived flows such as DNS and ICMP, instead of querying, the forwarder immediately sends the packet to the director, which then sends them to the owner. A connection can have multiple forwarders; the most efficient throughput is achieved by a good load-balancing method where there are no forwarders and all packets of a connection are received by the owner.

Note We do not recommend disabling TCP sequence randomization when using clustering. There is a small chance that some TCP sessions won't be established, because the SYN/ACK packet might be dropped.

- Fragment Owner—For fragmented packets, cluster nodes that receive a fragment determine a fragment owner using a hash of the fragment source IP address, destination IP address, and the packet ID. All fragments are then forwarded to the fragment owner over the cluster control link. Fragments may be load-balanced to different cluster nodes, because only the first fragment includes the 5-tuple used in the switch load balance hash. Other fragments do not contain the source and destination ports and may be load-balanced to other cluster nodes. The fragment owner temporarily reassembles the packet so it can determine the director based on a hash of the source/destination IP address and ports. If it is a new connection, the fragment owner will register to be the connection owner. If it is an existing connection, the fragment owner forwards all fragments to the provided connection owner over the cluster control link. The connection owner will then reassemble all fragments.

Port Address Translation Connections

New Connection Ownership

Traffic redirection is not supported in this release. When a new connection is directed to a node of the cluster via load balancing, that node owns both directions of the connection. All the subsequent packets for the same connection should arrive the same node. If any connection packets arrive at a different node, they will be dropped. If a reverse flow arrives at a different node, it will be dropped as well. For centralized features, if the connections do not arrive on the control node, they will be dropped.

By default, AWS GWLB uses 5-tuple to maintain flow stickiness. It is recommended to enable 2-tuple or 3-tuple stickiness on AWS GWLB to ensure the same flows are sent to the same node.

Sample Data Flow for TCP

The following example shows the establishment of a new connection.

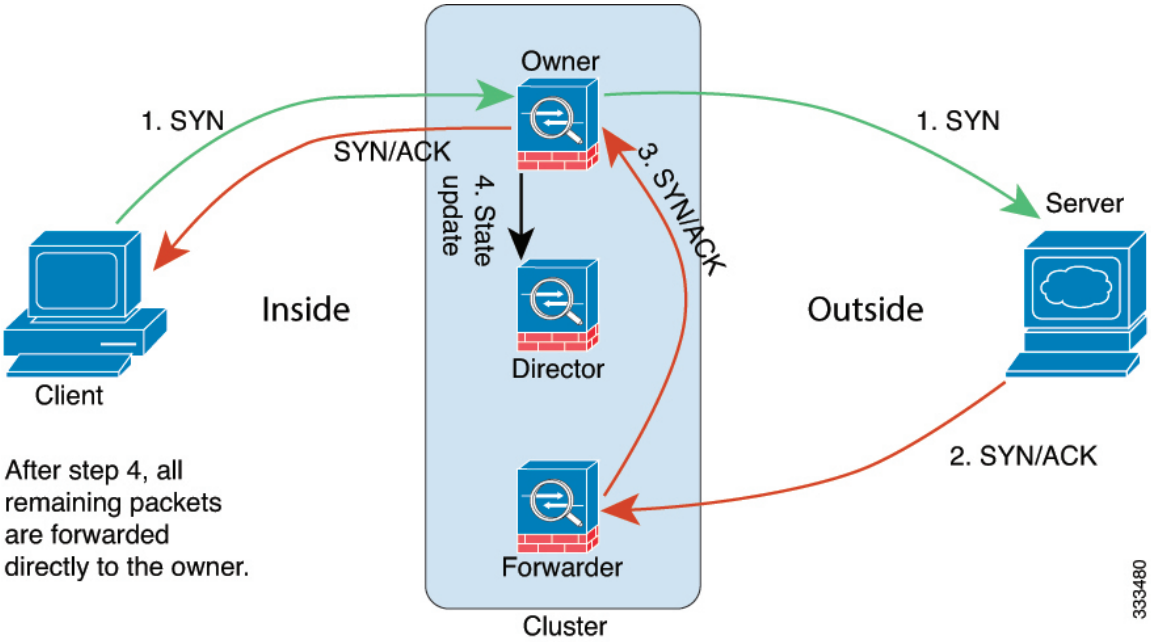

1. The SYN packet originates from the client and is delivered to one Firewall Threat Defense (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.

2. The SYN-ACK packet originates from the server and is delivered to a different Firewall Threat Defense (based on the load balancing method). This Firewall Threat Defense is the forwarder.
3. Because the forwarder does not own the connection, it decodes owner information from the SYN cookie, creates a forwarding flow to the owner, and forwards the SYN-ACK to the owner.
4. The owner sends a state update to the director, and forwards the SYN-ACK to the client.
5. The director receives the state update from the owner, creates a flow to the owner, and records the TCP state information as well as the owner. The director acts as the backup owner for the connection.
6. Any subsequent packets delivered to the forwarder will be forwarded to the owner.
7. If packets are delivered to any additional nodes, it will query the director for the owner and establish a flow.
8. Any state change for the flow results in a state update from the owner to the director.

Sample Data Flow for ICMP and UDP

The following example shows the establishment of a new connection.

1. ***Figure 233: ICMP and UDP Data Flow***

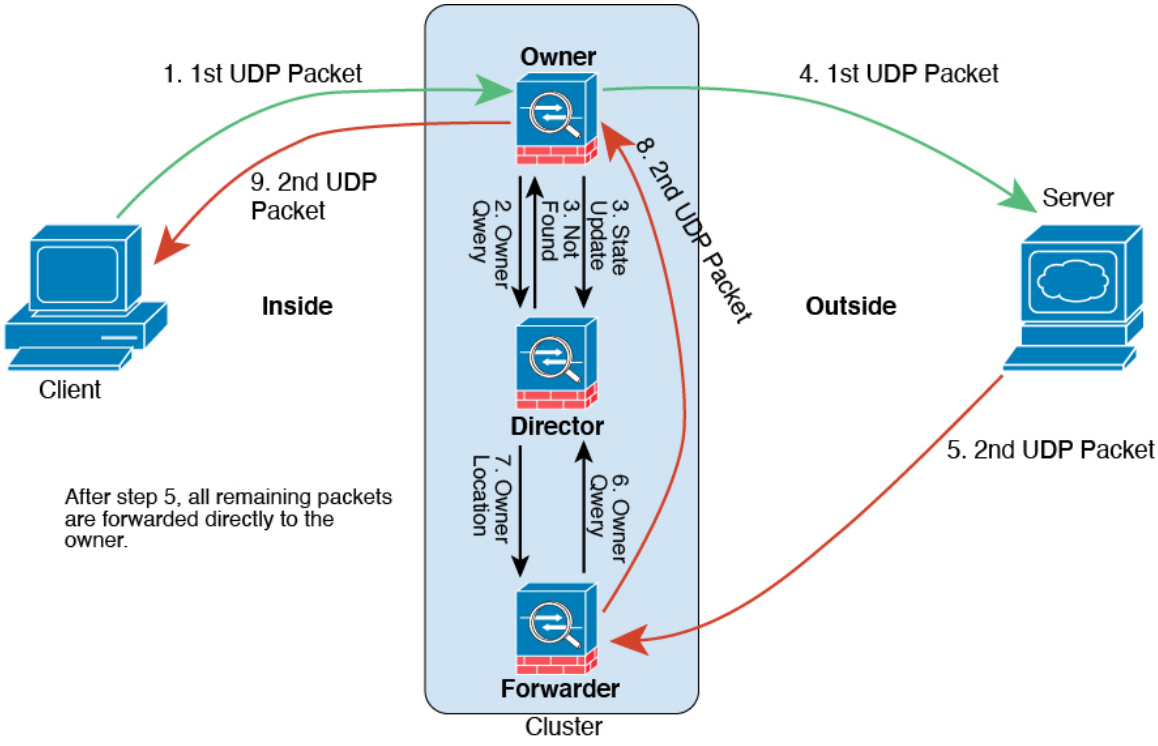

The first UDP packet originates from the client and is delivered to one Firewall Threat Defense (based on the load balancing method).

2. The node that received the first packet queries the director node that is chosen based on a hash of the source/destination IP address and ports.

3. The director finds no existing flow, creates a director flow and forwards the packet back to the previous node. In other words, the director has elected an owner for this flow.
4. The owner creates the flow, sends a state update to the director, and forwards the packet to the server.
5. The second UDP packet originates from the server and is delivered to the forwarder.
6. The forwarder queries the director for ownership information. For short-lived flows such as DNS, instead of querying, the forwarder immediately sends the packet to the director, which then sends it to the owner.
7. The director replies to the forwarder with ownership information.
8. The forwarder creates a forwarding flow to record owner information and forwards the packet to the owner.
9. The owner forwards the packet to the client.

History for Clustering

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
16-node clusters for the Secure Firewall 3100/4200.	7.6.0	7.6.0	For the Secure Firewall 3100 and 4200, the maximum nodes were increased from 8 to 16.
Individual interface mode for Secure Firewall 3100/4200 clusters.	7.6.0	7.6.0	Individual interfaces are normal routed interfaces, each with their own local IP address used for routing. The main cluster IP address for each interface is a fixed address that always belongs to the control node. When the control node changes, the main cluster IP address moves to the new control node, so management of the cluster continues seamlessly. Load balancing must be configured separately on the upstream switch. Restrictions: Not supported for container instances. New/modified screens: • **Devices** > **Device Management** > **Add Cluster** • **Devices** > **Device Management** > **Cluster** > **Interfaces / EIGRP / OSPF / OSPFv3 / BGP** • **Objects** > **Object Management** > **Address Pools** > **MAC Address Pool**
MTU ping test on cluster node join	7.6.0	7.6.0	When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the ping fails, a notification is generated so you can fix the MTU mismatch on connecting switches and try again.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Cluster control link ping tool.	7.2.6 7.4.1	Any	You can check to make sure all the cluster nodes can reach each other over the cluster control link by performing a ping. One major cause for the failure of a node to join the cluster is an incorrect cluster control link configuration; for example, the cluster control link MTU may be set higher than the connecting switch MTUs. New/modified screens: **Devices** > **Device Management** > **More** > **Cluster Live Status.** Version restrictions: Not supported with Firewall Management Center Version 7.3.x or 7.4.0.
Troubleshooting file generation and download available from Device and Cluster pages.	7.4.1	7.4.1	You can generate and download troubleshooting files for each device on the Device page and also for all cluster nodes on the Cluster page. For a cluster, you can download all files as a single compressed file. You can also include cluster logs for the cluster for cluster nodes. You can alternatively trigger file generation from the **Devices** > **Device Management** > **More** > **Troubleshoot Files** menu. New/modified screens: • **Devices** > **Device Management** > **Device** > **General** • **Devices** > **Device Management** > **Cluster** > **General**
Automatic generation of a troubleshooting file on a node when it fails to join the cluster.	7.4.1	7.4.1	If a node fails to join the cluster, a troubleshooting file is automatically generated for the node. You can download the file from **Tasks** or from the **Cluster** page.
View CLI output for a device or device cluster.	7.4.1	Any	You can view a set of pre-defined CLI outputs that can help you troubleshoot the device or cluster. You can also enter any **show** command and see the output. New/modified screens: **Devices** > **Device Management** > **Cluster** > **General**
Clustering for the Secure Firewall 4200	7.4.0	7.4.0	The Secure Firewall 4200 supports Spanned EtherChannel clustering for up to 8 nodes.
Cluster health monitor settings	7.3.0	Any	You can now edit cluster health monitor settings. New/modified screens: **Devices** > **Device Management** > **Cluster** > **Cluster Health Monitor Settings** **Note** If you previously configured these settings using FlexConfig, be sure to remove the FlexConfig configuration before you deploy. Otherwise the FlexConfig configuration will overwrite the management center configuration.
Cluster health monitor dashboard	7.3.0	Any	You can now view cluster health on the cluster health monitor dashboard. New/modified screens: **System** > **Health** > **Monitor**

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Automatic configuration of the cluster control link MTU	7.2.0	7.2.0	The MTU of the cluster control link interface is now automatically set to 100 bytes more than the highest data interface MTU; by default, the MTU is 1600 bytes.
Clustering for the Secure Firewall 3100	7.1.0	7.1.0	The Secure Firewall 3100 supports Spanned EtherChannel clustering for up to 8 nodes. New/modified screens: • **Devices** > **Device Management** > **Add Cluster** • **Devices** > **Device Management** > **More** menu • **Devices** > **Device Management** > **Cluster** Supported platforms: Secure Firewall 3100

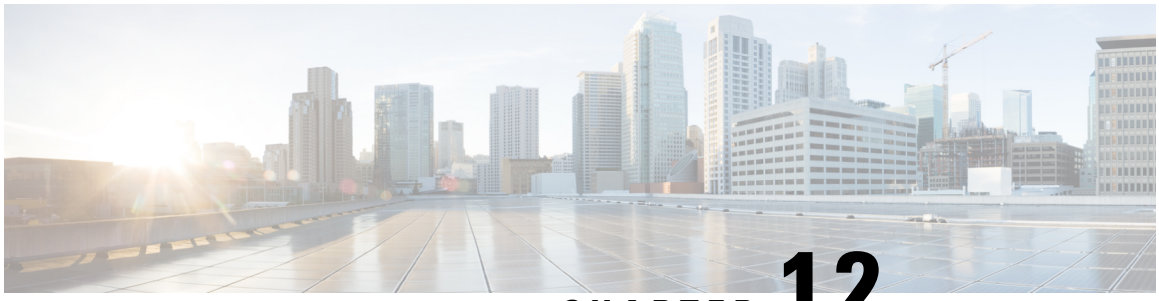

CHAPTER 12

Clustering for Threat Defense Virtual in a Private Cloud

Clustering lets you group multiple Firewall Threat Defense Virtuals together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices. You can deploy Firewall Threat Defense Virtual clusters in a private cloud using VMware and KVM. Only routed firewall mode is supported.

Note Some features are not supported when using clustering. See Unsupported Features and Clustering, on page 564.

- About Threat Defense Virtual Clustering in the Private Cloud, on page 525
- Licenses for Threat Defense Virtual Clustering, on page 529
- Requirements and Prerequisites for Threat Defense Virtual Clustering, on page 530
- Guidelines for Threat Defense Virtual Clustering, on page 531
- Configure Threat Defense Virtual Clustering, on page 532
- Manage Cluster Nodes, on page 546
- Monitoring the Cluster, on page 556
- Troubleshooting the Cluster, on page 562
- Reference for Clustering, on page 564
- History for Threat Defense Virtual Clustering in a Private Cloud, on page 576

About Threat Defense Virtual Clustering in the Private Cloud

This section describes the clustering architecture and how it works.

How the Cluster Fits into Your Network

The cluster consists of multiple firewalls acting as a single device. To act as a cluster, the firewalls need the following infrastructure:

- Isolated network for intra-cluster communication, known as the *cluster control link*, using VXLAN interfaces. VXLANs, which act as Layer 2 virtual networks over Layer 3 physical networks, let the Firewall Threat Defense Virtual send broadcast/multicast messages over the cluster control link.

- Management access to each firewall for configuration and monitoring. The Firewall Threat Defense Virtual deployment includes a Management 0/0 interface that you will use to manage the cluster nodes.

When you place the cluster in your network, the upstream and downstream routers need to be able to load-balance the data coming to and from the cluster using Layer 3 Individual interfaces and one of the following methods:

- Policy-Based Routing—The upstream and downstream routers perform load balancing between nodes using route maps and ACLs.
- Equal-Cost Multi-Path Routing—The upstream and downstream routers perform load balancing between nodes using equal cost static or dynamic routes.

Note Layer 2 Spanned EtherChannels are not supported.

Control and Data Node Roles

One member of the cluster is the control node. If multiple cluster nodes come online at the same time, the control node is determined by the priority setting; the priority is set between 1 and 100, where 1 is the highest priority. All other members are data nodes. When you first create the cluster, you specify which node you want to be the control node, and it will become the control node simply because it is the first node added to the cluster.

All nodes in the cluster share the same configuration. The node that you initially specify as the control node will overwrite the configuration on the data nodes when they join the cluster, so you only need to perform initial configuration on the control node before you form the cluster.

Some features do not scale in a cluster, and the control node handles all traffic for those features.

Individual Interfaces

You can configure cluster interfaces as *Individual interfaces*.

Individual interfaces are normal routed interfaces, each with their own *Local IP address* used for routing. The *Main cluster IP address* for each interface is a fixed address that always belongs to the control node. When the control node changes, the Main cluster IP address moves to the new control node, so management of the cluster continues seamlessly.

IPS-only interfaces (inline sets and passive interfaces) are not supported as Individual interfaces.

Because interface configuration must be configured only on the control node, you configure a pool of IP addresses to be used for a given interface on the cluster nodes, including one for the control node.

Load balancing must be configured separately on the upstream switch.

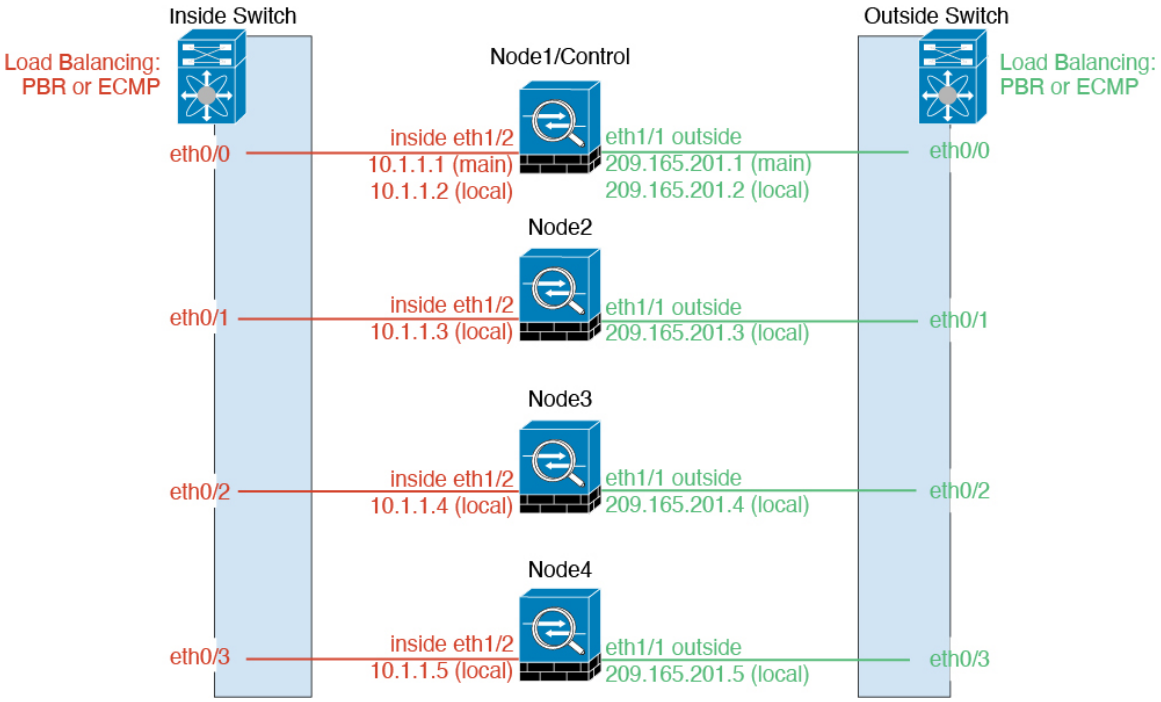

Note Layer 2 Spanned EtherChannels are not supported.

Policy-Based Routing

When using Individual interfaces, each Firewall Threat Defense interface maintains its own IP address and MAC address. One method of load balancing is Policy-Based Routing (PBR).

We recommend this method if you are already using PBR, and want to take advantage of your existing infrastructure.

PBR makes routing decisions based on a route map and ACL. You must manually divide traffic between all Firewall Threat Defenses in a cluster. Because PBR is static, it may not achieve the optimum load balancing result at all times. To achieve the best performance, we recommend that you configure the PBR policy so that forward and return packets of a connection are directed to the same Firewall Threat Defense. For example, if you have a Cisco router, redundancy can be achieved by using Cisco IOS PBR with Object Tracking. Cisco IOS Object Tracking monitors each Firewall Threat Defense using ICMP ping. PBR can then enable or disable route maps based on reachability of a particular Firewall Threat Defense. See the following URLs for more details:

http://www.cisco.com/c/en/us/solutions/data-center-virtualization/intelligent-traffic-director/index.html

http://www.cisco.com/en/US/products/ps6599/products_white_paper09186a00800a4409.shtml

Equal-Cost Multi-Path Routing

When using Individual interfaces, each Firewall Threat Defense interface maintains its own IP address and MAC address. One method of load balancing is Equal-Cost Multi-Path (ECMP) routing.

We recommend this method if you are already using ECMP, and want to take advantage of your existing infrastructure.

ECMP routing can forward packets over multiple “best paths” that tie for top place in the routing metric. Like EtherChannel, a hash of source and destination IP addresses and/or source and destination ports can be used to send a packet to one of the next hops. If you use static routes for ECMP routing, then the Firewall Threat Defense failure can cause problems; the route continues to be used, and traffic to the failed Firewall Threat Defense will be lost. If you use static routes, be sure to use a static route monitoring feature such as Object Tracking. We recommend using dynamic routing protocols to add and remove routes, in which case, you must configure each Firewall Threat Defense to participate in dynamic routing.

Cluster Control Link

Each node must dedicate one interface as a VXLAN (VTEP) interface for the cluster control link. For more information about VXLAN, see Configure VXLAN Interfaces, on page 848.

VXLAN Tunnel Endpoint

VXLAN tunnel endpoint (VTEP) devices perform VXLAN encapsulation and decapsulation. Each VTEP has two interface types: one or more virtual interfaces called VXLAN Network Identifier (VNI) interfaces, and a regular interface called the VTEP source interface that tunnels the VNI interfaces between VTEPs. The VTEP source interface is attached to the transport IP network for VTEP-to-VTEP communication.

VTEP Source Interface

The VTEP source interface is a regular Firewall Threat Defense Virtual interface with which you plan to associate the VNI interface. You can configure one VTEP source interface to act as the cluster control link. The source interface is reserved for cluster control link use only. Each VTEP source interface has an IP address on the same subnet. This subnet should be isolated from all other traffic, and should include only the cluster control link interfaces.

VNI Interface

A VNI interface is similar to a VLAN interface: it is a virtual interface that keeps network traffic separated on a given physical interface by using tagging. You can only configure one VNI interface. Each VNI interface has an IP address on the same subnet.

Peer VTEPs

Unlike regular VXLAN for data interfaces, which allows a single VTEP peer, The Firewall Threat Defense Virtual clustering allows you to configure multiple peers.

Cluster Control Link Traffic Overview

Cluster control link traffic includes both control and data traffic.

Control traffic includes:

- Control node election.
- Configuration replication.
- Health monitoring.

Data traffic includes:

- State replication.
- Connection ownership queries and data packet forwarding.

Configuration Replication

All nodes in the cluster share a single configuration. You can only make configuration changes on the control node (with the exception of the bootstrap configuration), and changes are automatically synced to all other nodes in the cluster.

Management Network

You must manage each node using the Management interface; management from a data interface is not supported with clustering.

Licenses for Threat Defense Virtual Clustering

Each Firewall Threat Defense Virtual cluster node requires the same performance tier license. We recommend using the same number of CPUs and memory for all members, or else performance will be limited on all nodes to match the least capable member. The throughput level will be replicated from the control node to each data node so they match.

You assign feature licenses to the cluster as a whole, not to individual nodes. However, each node of the cluster consumes a separate license for each feature. The clustering feature itself does not require any licenses.

When you add the control node to the Firewall Management Center, you can specify the feature licenses you want to use for the cluster. Before you create the cluster, it doesn't matter which licenses are assigned to the data nodes; the license settings for the control node are replicated to each of the data nodes. You can modify licenses for the cluster by clicking **Edit Licenses** in **System** (⚙) > **Licenses** > **Smart Licenses** or choosing**Devices** > **Device Management**, clicking **Edit** (✎) for the cluster, and then in the **Cluster** > **License** area, clicking **Edit** (✎).

Note If you add the cluster before the Firewall Management Center is licensed (and running in Evaluation mode), then when you license the Firewall Management Center, you can experience traffic disruption when you deploy policy changes to the cluster. Changing to licensed mode causes all data units to leave the cluster and then rejoin.

Requirements and Prerequisites for Threat Defense Virtual Clustering

Model Requirements

- FTDv5, FTDv10, FTDv20, FTDv30, FTDv50, FTDv100
- VMware or KVM
- A maximum of 16 nodes in a cluster in a 4x4 configuration is supported. You can set up a maximum of four hosts with a maximum of four threat defense virtual instances in each host.

User Roles

- Admin
- Access Admin
- Network Admin

Hardware and Software Requirements

All units in a cluster:

- Must have jumbo frame reservation enabled for the cluster control link. Do this in the Day 0 configuration when you deploy the Firewall Threat Defense Virtual by setting `"DeploymentType": "Cluster"`. Otherwise, you must restart each node to enable jumbo frames after the cluster has formed and is healthy.
- (KVM only) Must use CPU hard partitioning (CPU pinning) for all VMs on the KVM host.
- Must be the same performance tier. We recommend using the same number of CPUs and memory for all nodes, or performance will be limited on all nodes to match the least capable node.
- Must use the management interface for Firewall Management Center communications. Data interface management is not supported.
- Must run the same version, except during upgrade. Hitless upgrade is supported.
- Must be in the same domain.
- Must be in the same group.
- Must not have any deployment pending or in progress.
- Must not have any unsupported features configured on the control node: Unsupported Features and Clustering, on page 564.
- Must not have VPN configured on the data nodes. The control node can have site-to-site VPN configured.

Firewall Management Center Requirements

Make sure the Firewall Management Center NTP server is set to a reliable server that is reachable by all cluster nodes to ensure proper clock sync. By default, the device uses the same NTP server as the Firewall Management

Center. If the time is not set to be the same on all cluster nodes, they can be removed automatically from the cluster.

Switch Requirements

Be sure to complete the switch configuration before you configure clustering. Make sure the ports connected to the cluster control link have the correct (higher) MTU configured. By default, the cluster control link MTU is set to 154 bytes higher than the data interfaces. If the switches have an MTU mismatch, the cluster formation will fail.

Guidelines for Threat Defense Virtual Clustering

High Availability

High Availability is not supported with clustering.

IPv6

The cluster control link is only supported using IPv4.

Additional Guidelines

- When significant topology changes occur (such as adding or removing an EtherChannel interface, enabling or disabling an interface on the Firewall Threat Defense Virtual, adding an additional switch to form a VSS or vPC, configuring IP addresses or interface flap on the cluster) you should disable the health check feature and also disable interface monitoring for the interfaces that are affected by the topology changes. When the topology change is complete, and the configuration change is synced to all units, you can re-enable the interface health check feature.
- When adding a unit to an existing cluster, or when reloading a unit, there will be a temporary, limited packet/connection drop; this is expected behavior. In some cases, the dropped packets can hang your connection; for example, dropping a FIN/ACK packet for an FTP connection will make the FTP client hang. In this case, you need to reestablish the FTP connection.
- For decrypted TLS/SSL connections, the decryption states are not synchronized, and if the connection owner fails, then decrypted connections will be reset. New connections will need to be established to a new unit. Connections that are not decrypted (they match a do-not-decrypt rule) are not affected and are replicated correctly.
- We do not support VXLANs for data interfaces; only the cluster control link supports VXLAN.

Defaults for Clustering

- The cLACP system ID is auto-generated, and the system priority is 1 by default.
- The cluster health check feature is enabled by default with the holdtime of 3 seconds. Interface health monitoring is enabled on all interfaces by default.
- The cluster auto-rejoin feature for a failed cluster control link is unlimited attempts every 5 minutes.
- The cluster auto-rejoin feature for a failed data interface is 3 attempts every 5 minutes, with the increasing interval set to 2.

- Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

Configure Threat Defense Virtual Clustering

To configure clustering after you deploy your Firewall Threat Defense Virtuals, perform the following tasks.

Add Nodes to the Management Center

Before configuring clustering, deploy each cluster node, then add the devices as standalone units on the Firewall Management Center.

Procedure

Step 1 Deploy each cluster node according the Cisco Secure Firewall Threat Defense Virtual Getting Started Guide.

All units in a cluster:

- Must have jumbo frame reservation enabled for the cluster control link. Do this in the Day 0 configuration when you deploy the Firewall Threat Defense Virtual by setting `"DeploymentType": "Cluster"`. Otherwise, you must restart each node to enable jumbo frames after the cluster has formed and is healthy.
- (KVM only) Must use CPU hard partitioning (CPU pinning) for all VMs on the KVM host.

Step 2 Add each node to the Firewall Management Center as a standalone device in the same domain and group.

See Add a Device Using a Registration Key—Basic Configuration, on page 31. You can create a cluster with a single device, and then add more nodes later. The initial settings (licensing, access control policy) that you set when you add a device will be inherited by all cluster nodes from the control node. You will choose the control node when forming the cluster.

Create a Cluster

Form a cluster from one or more devices in the Firewall Management Center.

Before you begin

Some features are not compatible with clustering, so you should wait to perform configuration until after you enable clustering. Some features will block cluster creation if they are already configured. For example, do not configure any IP addresses on interfaces, or unsupported interface types such as BVIs.

Procedure

Step 1 Choose **Devices** > **Device Management**, and then choose **Add** > **Cluster**.

The **Add Cluster Wizard** appears.

Figure 234: Add Cluster Wizard

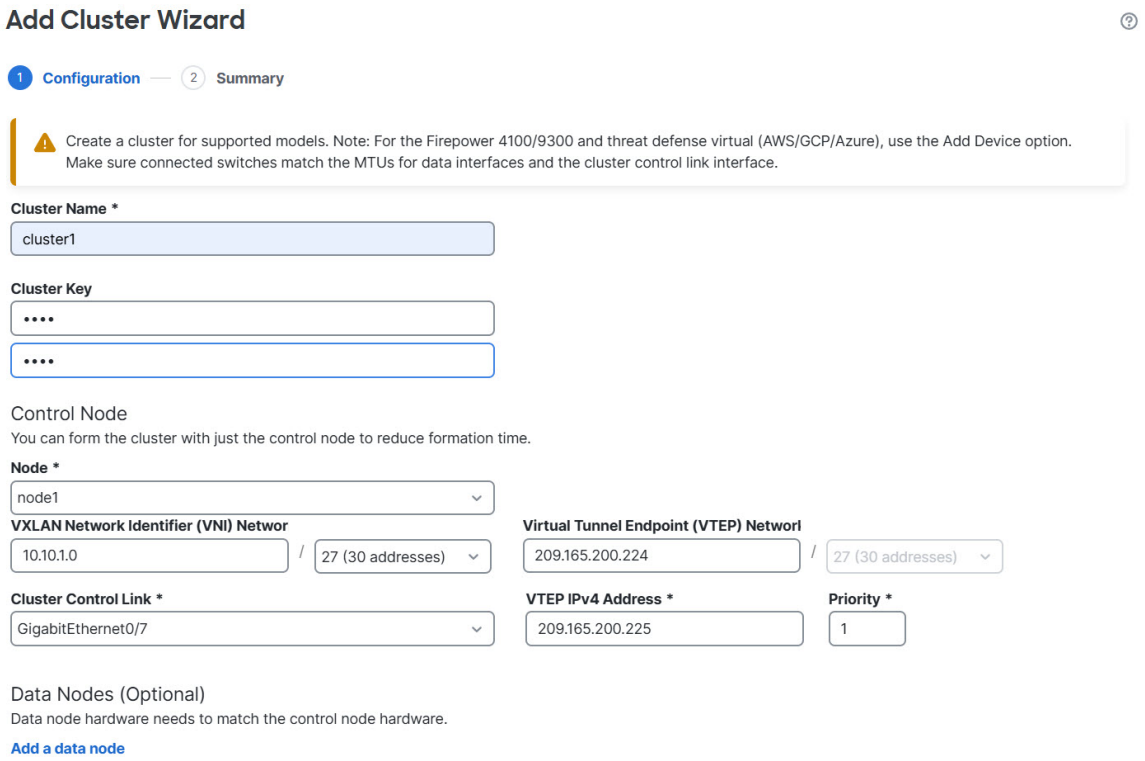

Step 2 Specify a **Cluster Name** and an authentication **Cluster Key** for control traffic.

- **Cluster Name**—An ASCII string from 1 to 38 characters.
- **Cluster Key**—An ASCII string from 1 to 63 characters. The **Cluster Key** value is used to generate the encryption key. This encryption does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear.

Step 3 For the **Control Node**, set the following:

- **Node**—Choose the device that you want to be the control node initially. When the Firewall Management Center forms the cluster, it will add this node to the cluster first so it will be the control node.

  Note

  If you see an **Error** ( ) icon next to the node name, click the icon to view configuration issues. You must cancel cluster formation, resolve the issues, and then return to cluster formation. For example:

Figure 235: Configuration Issues

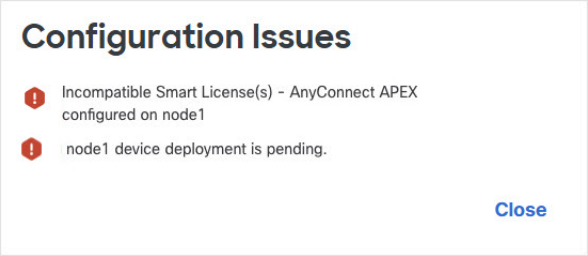

To resolve the above issues, remove the unsupported VPN license and deploy pending configuration changes to the device.

- **VXLAN Network Identifier (VNI) Network**—Specify an IPv4 subnet for the VNI network; IPv6 is not supported for this network. Specify a **24**, **25**, **26**, or **27** subnet. An IP address will be auto-assigned to each node on this network. The VNI network is the encrypted virtual network that runs on top of the physical VTEP network.
- **Cluster Control Link**—Choose the physical interface you want to use for the cluster control link.
- **Virtual Tunnel Endpoint (VTEP) Network**—Specify an IPv4 subnet for the physical interface network; IPv6 is not supported for this network. The VTEP network is a different network than the VNI network, and it is used for the physical cluster control link.
- **VTEP IPv4 Address**—This field will be auto-populated with the first address on the VTEP network.
- **Priority**—Set the priority of this node for control node elections. The priority is between 1 and 100, where 1 is the highest priority. Even if you set the priority to be lower than other nodes, this node will still be the control node when the cluster is first formed.

Step 4 For **Data Nodes (Optional)**, click **Add a data node** to add a node to the cluster.

You can form the cluster with only the control node for faster cluster formation, or you can add all nodes now. Set the following for each data node:

- **Node**—Choose the device that you want to add.

  Note

  If you see an **Error** ( ) icon next to the node name, click the icon to view configuration issues. You must cancel cluster formation, resolve the issues, and then return to cluster formation.

- **VTEP IPv4 Address**—This field will be auto-populated with the next address on the VTEP network.
- **Priority**—Set the priority of this node for control node elections. The priority is between 1 and 100, where 1 is the highest priority.

Step 5 Click **Continue**. Review the **Summary**, and then click **Save**.

The cluster bootstrap configuration is saved to the cluster nodes. The bootstrap configuration includes the VXLAN interface used for the cluster control link.

The cluster name shows on the **Devices** > **Device Management** page; expand the cluster to see the cluster nodes.

Figure 236: Cluster Management

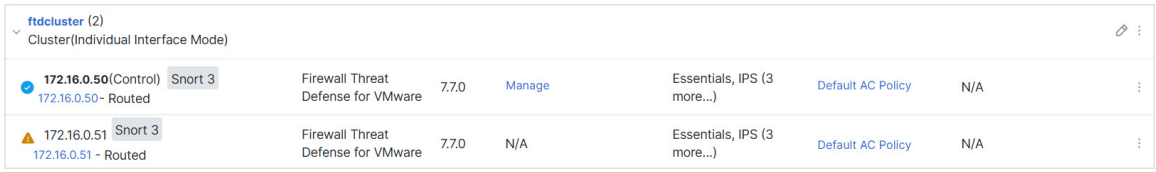

A node that is currently registering shows the loading icon.

Figure 237: Node Registration

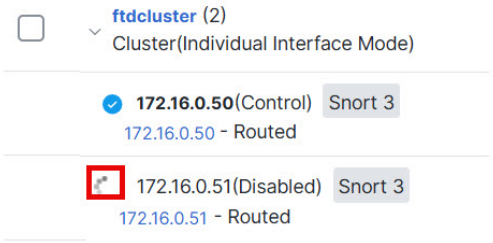

You can monitor cluster node registration by clicking the **Notifications** icon and choosing **Tasks**. The Firewall Management Center updates the Cluster Registration task as each node registers.

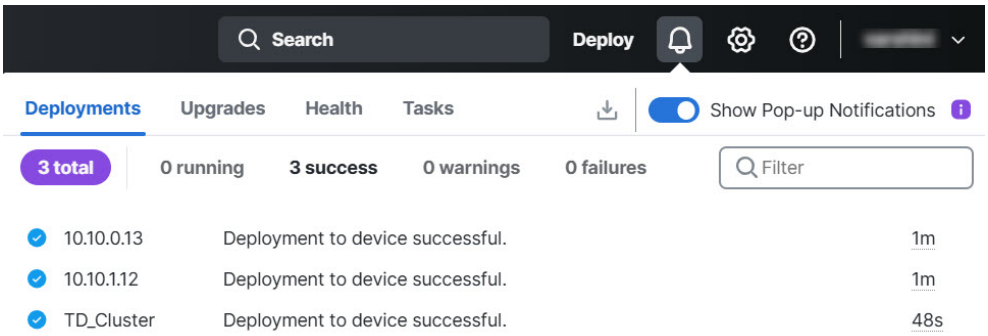

Step 6 Configure device-specific settings by clicking the **Edit** (✎) for the cluster.

Most configuration can be applied to the cluster as a whole, and not nodes in the cluster. For example, you can change the display name per node, but you can only configure interfaces for the whole cluster.

Step 7 On the **Devices** > **Device Management** > **Cluster** screen, you see **General** and other settings for the cluster.

Figure 238: Cluster Settings

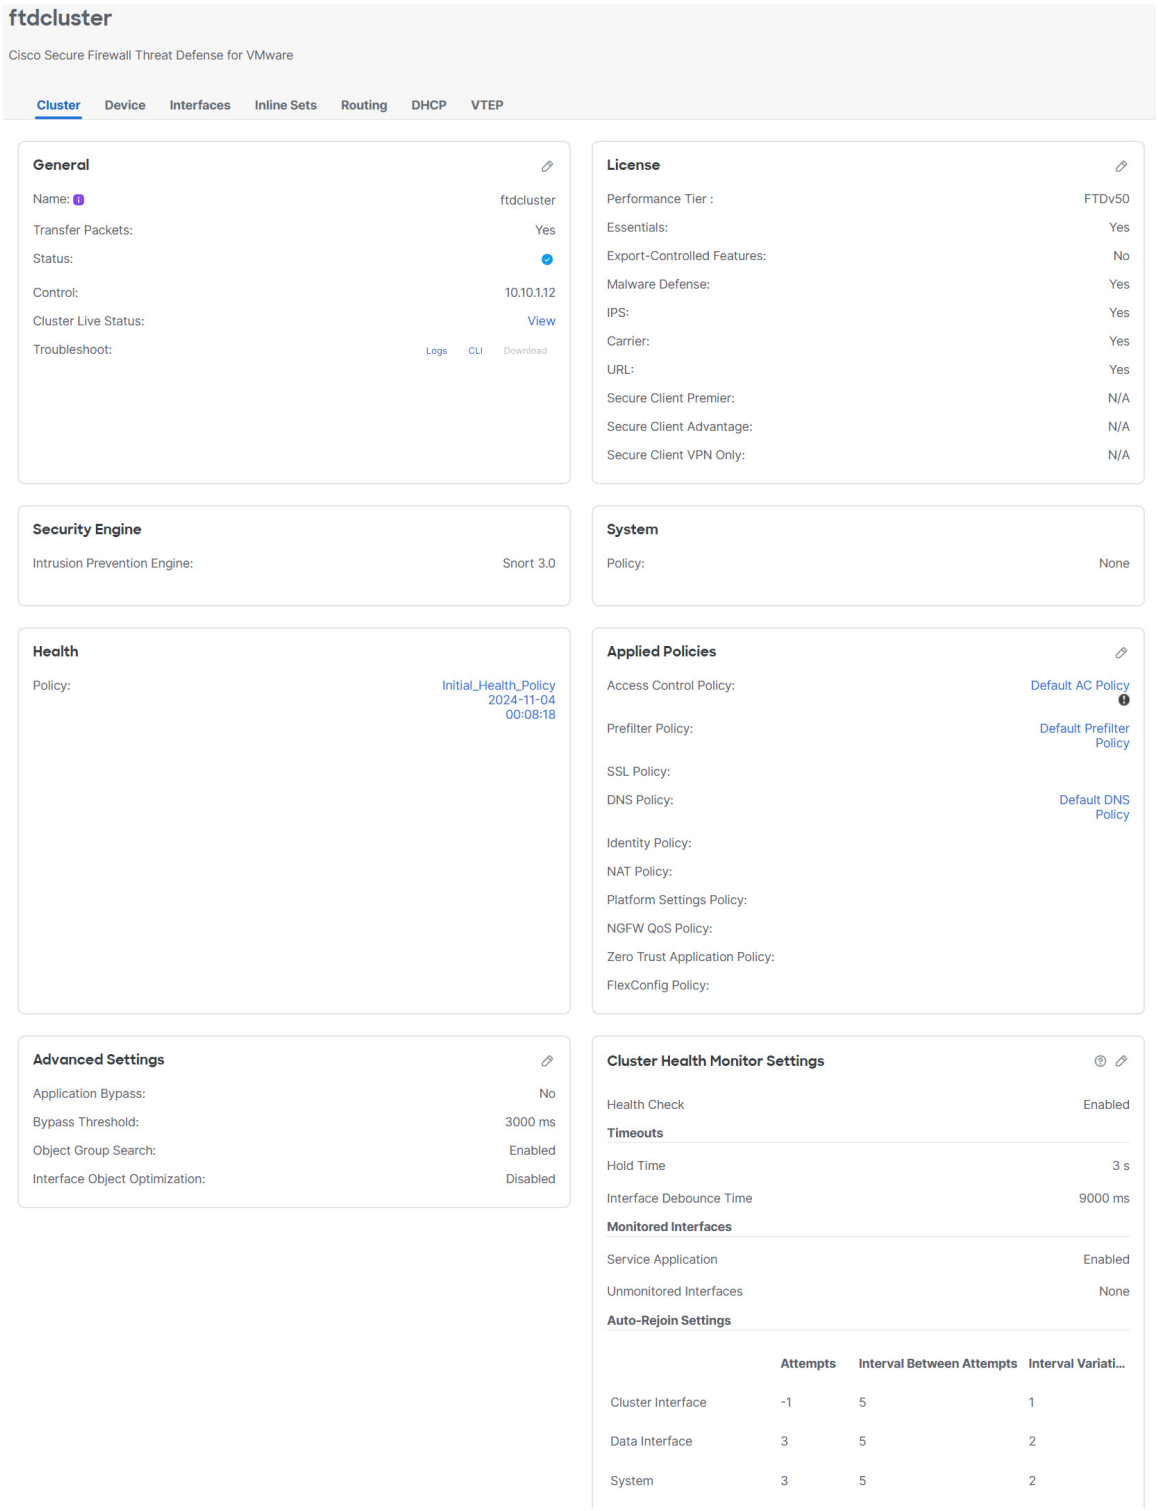

See the following cluster-specific items in the **General** area:

- **General > Name**—Change the cluster display name by clicking the **Edit** ( ).

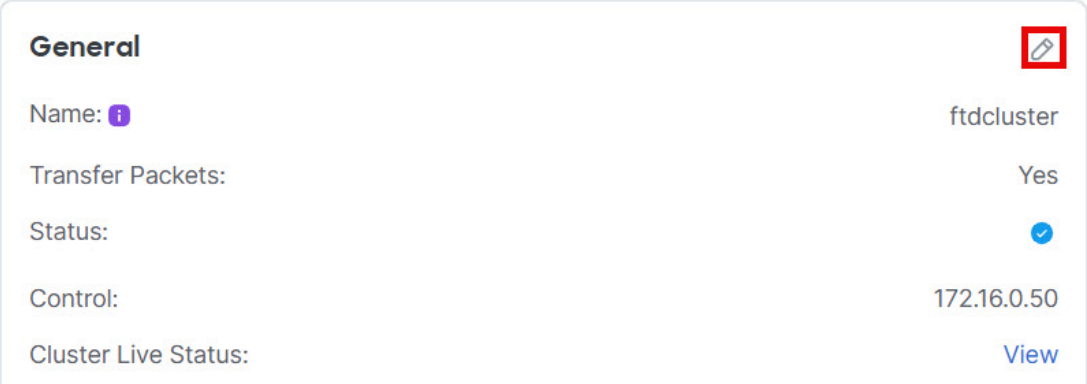

Then set the **Name** field.

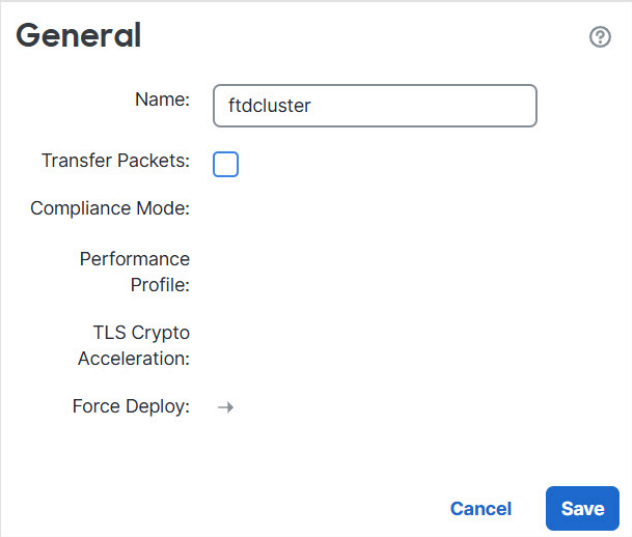

- **General > View**—Click the **View** link to open the **Cluster Status** dialog box.

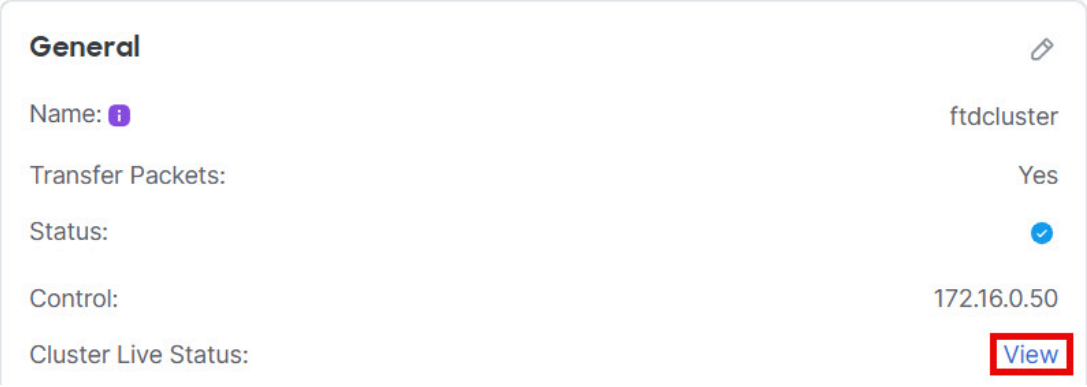

The **Cluster Status** dialog box also lets you retry data unit registration by clicking **Reconcile All**.You can also ping the cluster control link from a node. See Perform a Ping on the Cluster Control Link, on page 562.

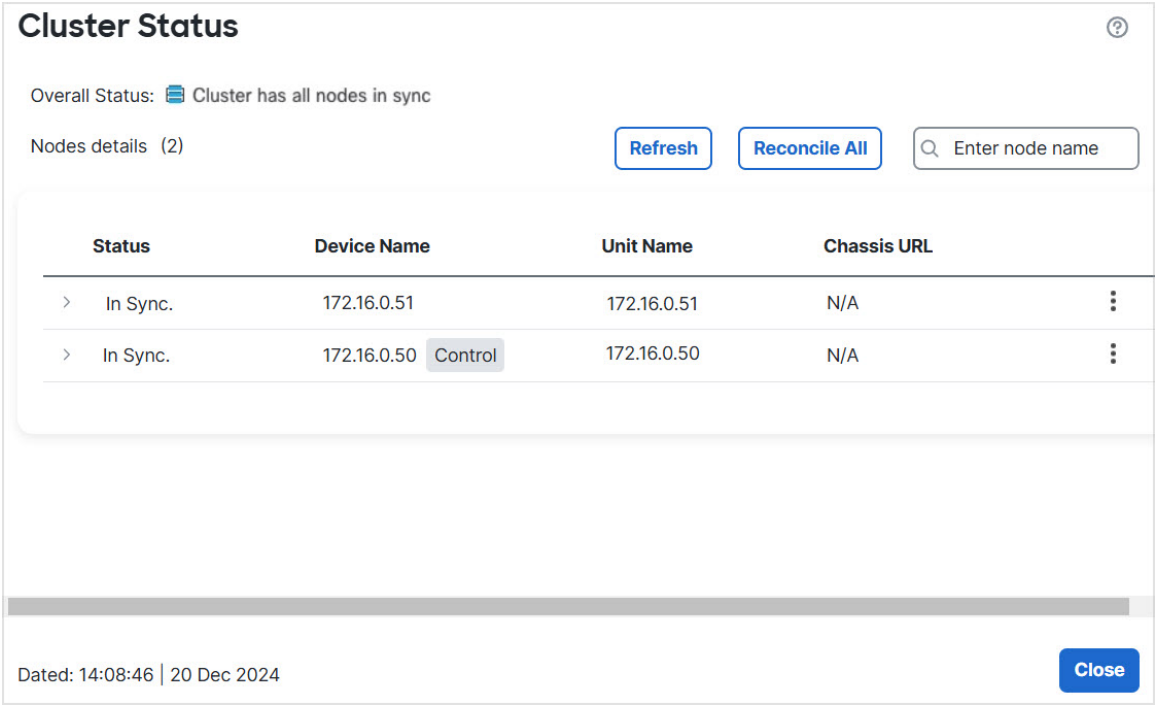

- **General > Troubleshoot**—You can generate and download troubleshooting logs, and you can view cluster CLIs. See Troubleshooting the Cluster, on page 562.

Figure 239: Troubleshoot

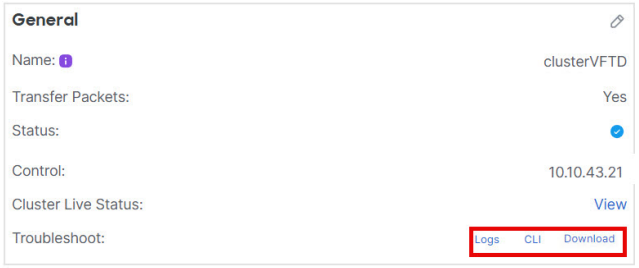

Step 8 On the **Devices** > **Device Management** > **Devices**, you can choose each member in the cluster from the top right drop-down menu and configure the following settings.

Figure 240: Device Settings

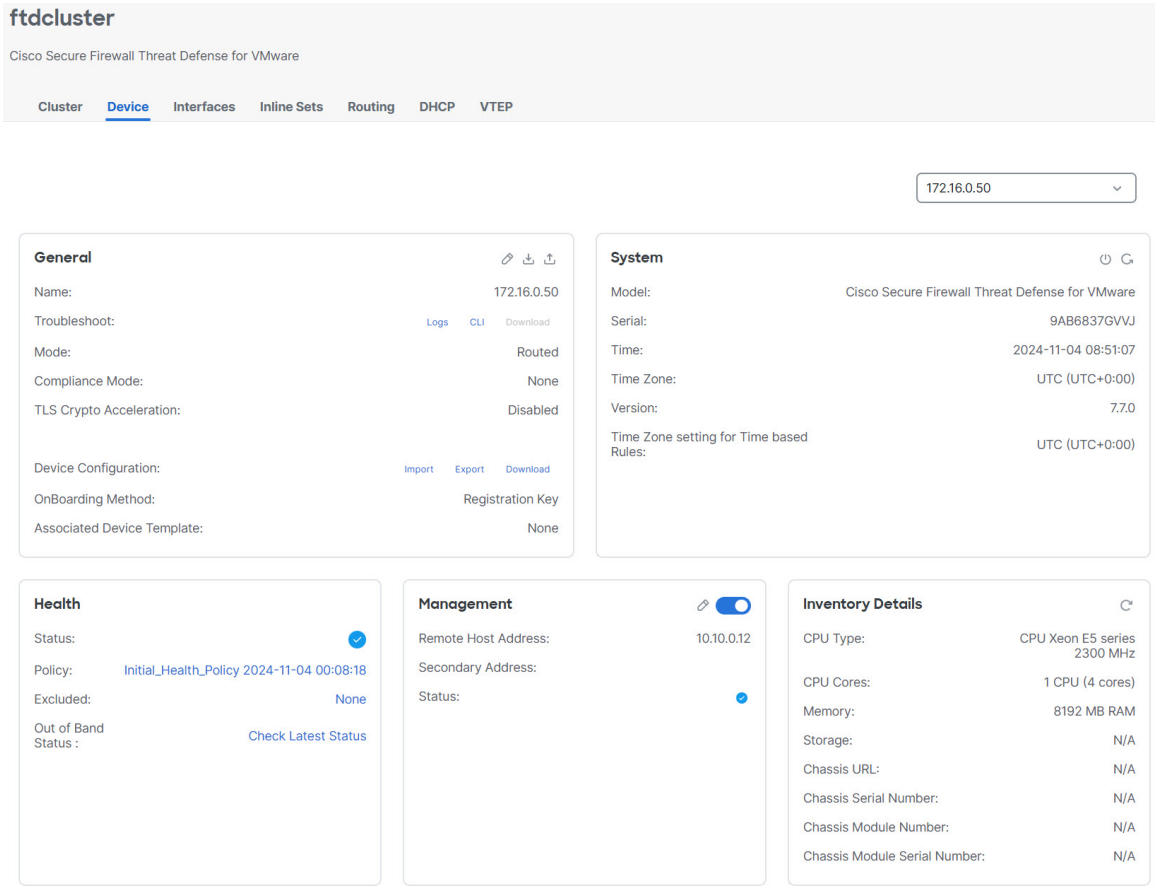

Figure 241: Choose Node

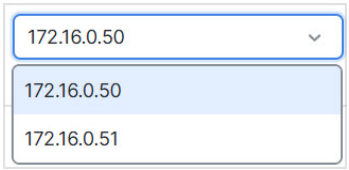

- **General > Name**—Change the cluster member display name by clicking the **Edit** (✎).

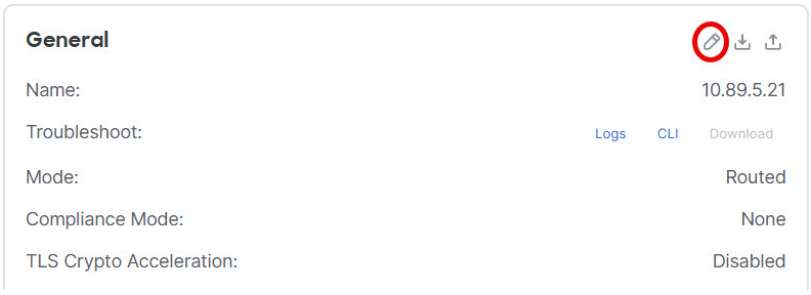

Then set the **Name** field.

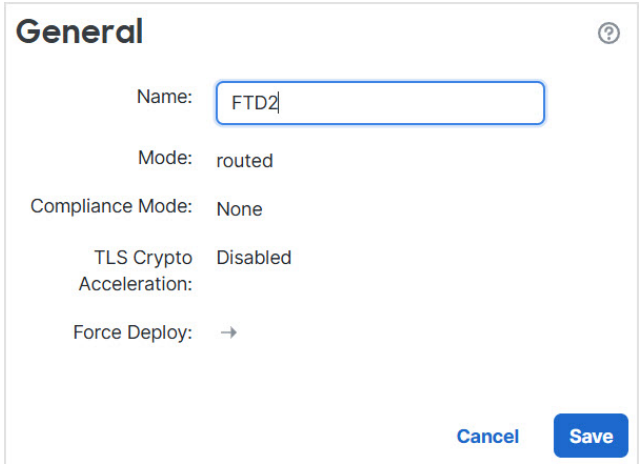

- **Management > Host**—If you change the management IP address in the device configuration, you must match the new address in the Firewall Management Center so that it can reach the device on the network. First disable the connection, edit the **Host** address in the **Management** area, then re-enable the connection.

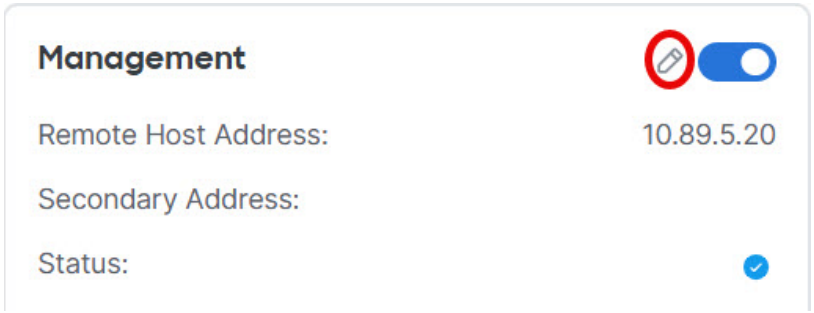

Step 9 If you deployed your cluster nodes without enabling jumbo-frame reservation, then restart all cluster nodes to enable jumbo frames, which are required for the cluster control link. See Shut Down or Restart the Device, on page 62.

If you previously enabled jumbo-frame reservation, you can skip this step.

Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead (100 bytes) and VXLAN overhead (54 bytes). When you create the cluster, the MTU is set to 154 bytes higher than the highest data interface MTU (1654 by default). If you later increase the data interface MTU, be sure to also increase the cluster control link MTU. For example, because the maximum MTU is 9198 bytes, then the highest data interface MTU can be 9044, while the cluster control link can be set to 9198. See Configure the MTU, on page 880.

Note
Make sure you configure switches connected to the cluster control link to the correct (higher) MTU; otherwise, cluster formation will fail. When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the ping fails, a notification is generated so you can fix the MTU mismatch on connecting switches and try again.

Configure Interfaces

This section describes how to configure interfaces to be Individual interfaces compatible with clustering. Individual interfaces are normal routed interfaces, each with their own IP address taken from a pool of IP addresses. The Main cluster IP address is a fixed address for the cluster that always belongs to the current control node. All data interfaces must be Individual interfaces.

Note You cannot use subinterfaces.

Procedure

Step 1 Choose **Objects** > **Object Management** > **Address Pools** to add an IPv4 and/or IPv6 address pool. See Address Pools, on page 1374.

Include at least as many addresses as there are units in the cluster. The Virtual IP address is not a part of this pool, but needs to be on the same network. You cannot determine the exact Local address assigned to each unit in advance.

Step 2 Choose **Devices** > **Device Management**, and click **Edit** (✎) next to the cluster.

Step 3 Click **Interfaces**, and then click **Edit** (✎) for a data interface.

Step 4 On the **IPv4**, enter the **IP Address** and mask. This IP address is a fixed address for the cluster, and always belongs to the current control unit.

Step 5 From the **IPv4 Address Pool** drop-down list, choose the address pool you created.

Note
If you want to manually assign a MAC address to this interface, you need to create a **mac-address pool** using FlexConfig.

Step 6 On **IPv6 > Basic**, from the **IPv6 Address Pool** drop-down list, choose the address pool you created.

Step 7 Configure other interface settings as normal.

Step 8 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Cluster Health Monitor Settings

The **Cluster Health Monitor Settings** section of the **Cluster** page displays the settings described in the table below.

Figure 242: Cluster Health Monitor Settings

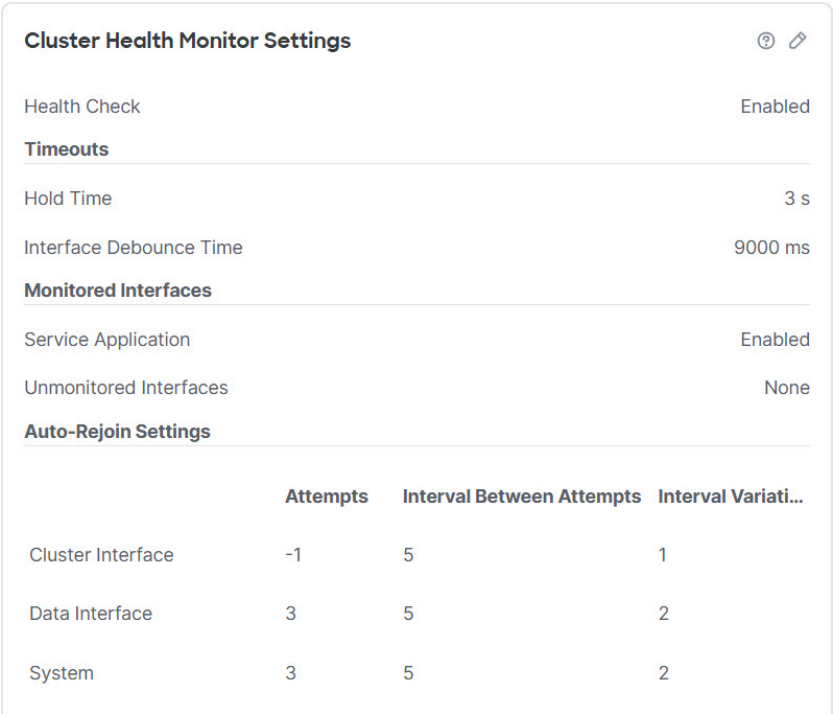

Table 34: Cluster Health Monitor Settings Section Table Fields

Field	Description
Timeouts	
Hold Time	Between .3 and 45 seconds; The default is 3 seconds. To determine node system health, the cluster nodes send heartbeat messages on the cluster control link to other nodes. If a node does not receive any heartbeat messages from a peer node within the hold time period, the peer node is considered unresponsive or dead.
Interface Debounce Time	Between 300 and 9000 ms. The default is 500 ms. The interface debounce time is the amount of time before the node considers an interface to be failed, and the node is removed from the cluster.
Monitored Interfaces	The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The amount of time before the node removes a member from the cluster depends on the type of interface and whether the node is an established node or is joining the cluster.
Service Application	Shows whether the Snort and disk-full processes are monitored.
Unmonitored Interfaces	Shows unmonitored interfaces.
Auto-Rejoin Settings	

Field	Description
Cluster Interface	Shows the auto-rejoin settings after a cluster control link failure.
Attempts	Between -1 and 65535. The default is -1 (unlimited). Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 1x the interval duration. Defines if the interval duration increases at each attempt.
Data Interfaces	Shows the auto-rejoin settings after a data interface failure.
Attempts	Between -1 and 65535. The default is 3. Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 2x the interval duration. Defines if the interval duration increases at each attempt.
System	Shows the auto-rejoin settings after internal errors. Internal failures include: application sync timeout; inconsistent application statuses; and so on.
Attempts	Between -1 and 65535. The default is 3. Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 2x the interval duration. Defines if the interval duration increases at each attempt.

Note If you disable the system health check, fields that do not apply when the system health check is disabled will not show.

You can change these settings from this section.

You can monitor any port-channel ID, single physical interface ID, as well as the Snort and disk-full processes. Health monitoring is not performed on VLAN subinterfaces or virtual interfaces such as VNIs or BVIs. You cannot configure monitoring for the cluster control link; it is always monitored.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the cluster you want to modify, click **Edit** (✎).

Step 3 Click **Cluster**.

Step 4 In the **Cluster Health Monitor Settings** section, click **Edit** (✎).

Step 5 Disable the system health check by clicking the **Health Check** slider .

Figure 243: Disable the System Health Check

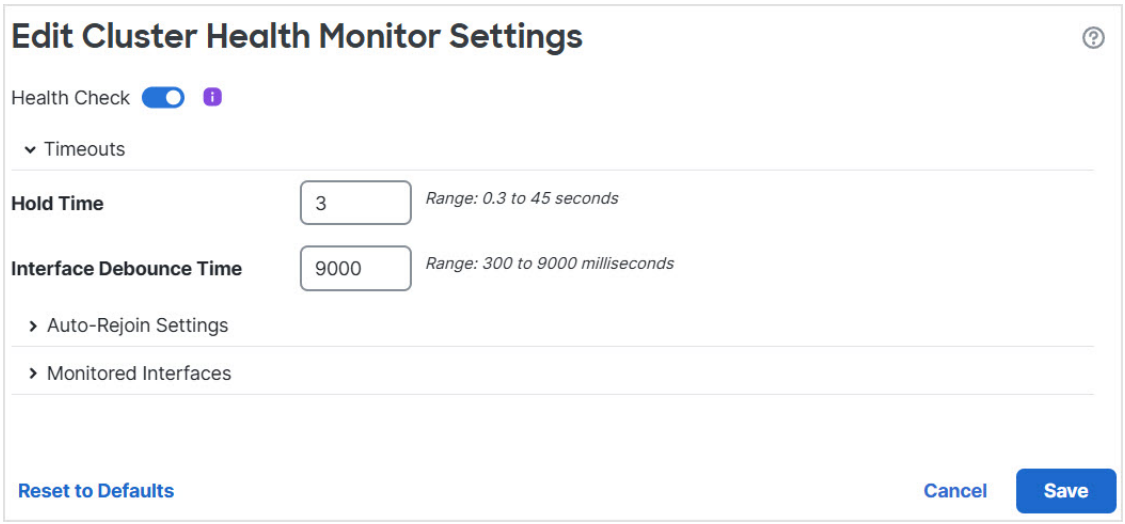

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 6 Configure the hold time and interface debounce time.

- **Hold Time**—Set the hold time to determine the amount of time between node heartbeat status messages, between .3 and 45 seconds; The default is 3 seconds.
- **Interface Debounce Time**—Set the debounce time between 300 and 9000 ms. The default is 500 ms. Lower values allow for faster detection of interface failures. Note that configuring a lower debounce time increases the chances of false-positives. When an interface status update occurs, the node waits the number of milliseconds specified before marking the interface as failed, and the node is removed from the cluster. In the case of an EtherChannel that transitions from a down state to an up state (for example, the switch reloaded, or the switch enabled an EtherChannel), a longer debounce time can prevent the interface from appearing to be failed on a cluster node just because another cluster node was faster at bundling the ports.

Step 7 Customize the auto-rejoin cluster settings after a health check failure.

Figure 244: Configure Auto-Rejoin Settings

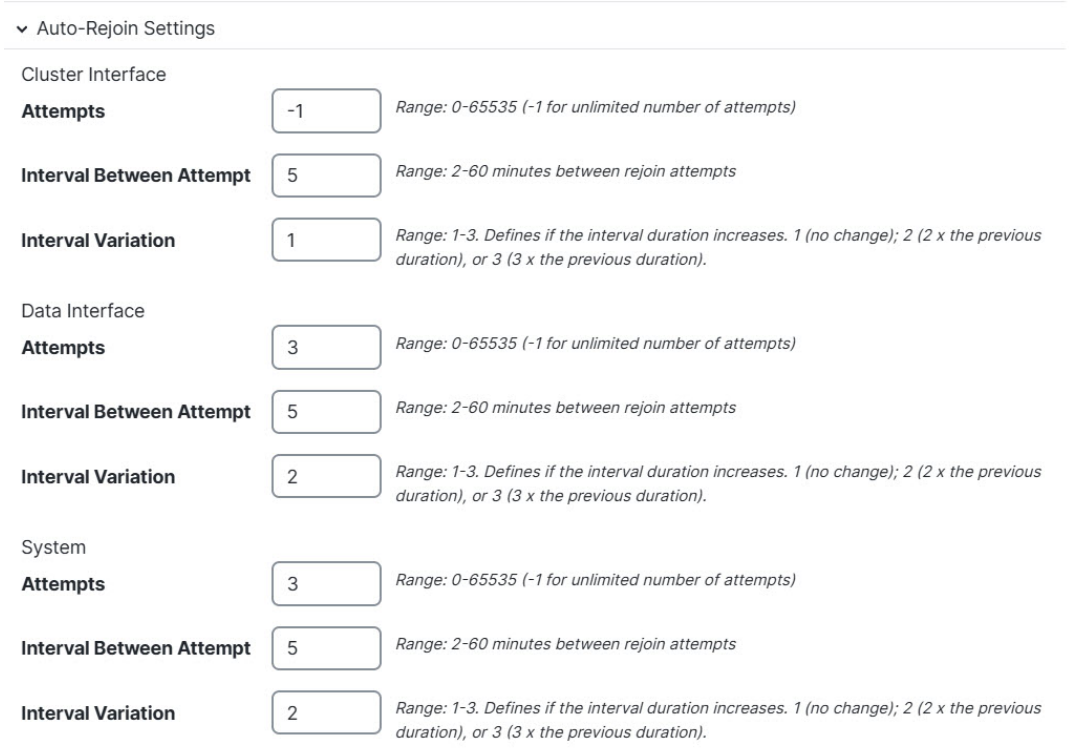

Set the following values for the **Cluster Interface**, **Data Interface**, and **System** (internal failures include: application sync timeout; inconsistent application statuses; and so on):

- **Attempts**—Sets the number of rejoin attempts, between -1 and 65535. **0** disables auto-rejoining. The default for the **Cluster Interface** is -1 (unlimited). The default for the **Data Interface** and **System** is 3.
- **Interval Between Attempts**—Defines the interval duration in minutes between rejoin attempts, between 2 and 60. The default value is 5 minutes. The maximum total time that the node attempts to rejoin the cluster is limited to 14400 minutes (10 days) from the time of last failure.
- **Interval Variation**—Defines if the interval duration increases. Set the value between 1 and 3: **1** (no change); **2** (2 x the previous duration), or **3** (3 x the previous duration). For example, if you set the interval duration to 5 minutes, and set the variation to 2, then the first attempt is after 5 minutes; the 2nd attempt is 10 minutes (2 x 5); the 3rd attempt 20 minutes (2 x 10), and so on. The default value is **1** for the **Cluster Interface** and **2** for the **Data Interface** and **System**.

Step 8 Configure monitored interfaces by moving interfaces in the **Monitored Interfaces** or **Unmonitored Interfaces** window. You can also check or uncheck **Enable Service Application Monitoring** to enable or disable monitoring of the Snort and disk-full processes.

Figure 245: Configure Monitored Interfaces

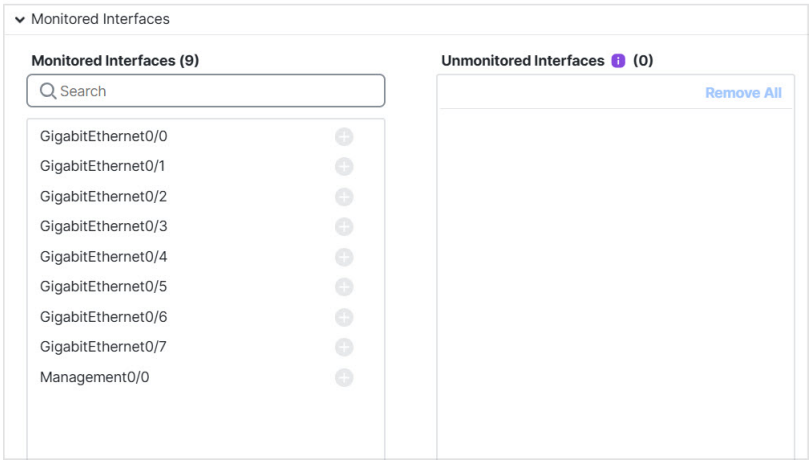

The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The amount of time before the node removes a member from the cluster depends on the type of interface and whether the node is an established node or is joining the cluster. Health check is enabled by default for all interfaces and for the Snort and disk-full processes.

You might want to disable health monitoring of non-essential interfaces.

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 9 Click **Save**.

Step 10 Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Manage Cluster Nodes

Add a New Cluster Node

You can add one or more new cluster nodes to an existing cluster.

Procedure

Step 1 Choose **Devices** > **Device Management**, click the **More** (⋮) for the cluster, and choose **Add Nodes**.

Figure 246: Add Nodes

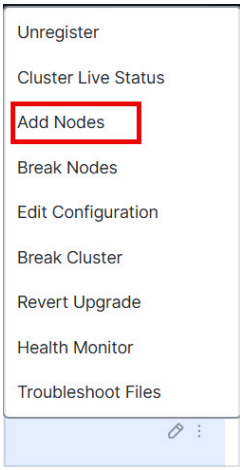

The **Manage Cluster Wizard** appears.

Step 2 From the **Node** menu, choose a device, and adjust the IP address and priority if desired.

Figure 247: Manage Cluster Wizard

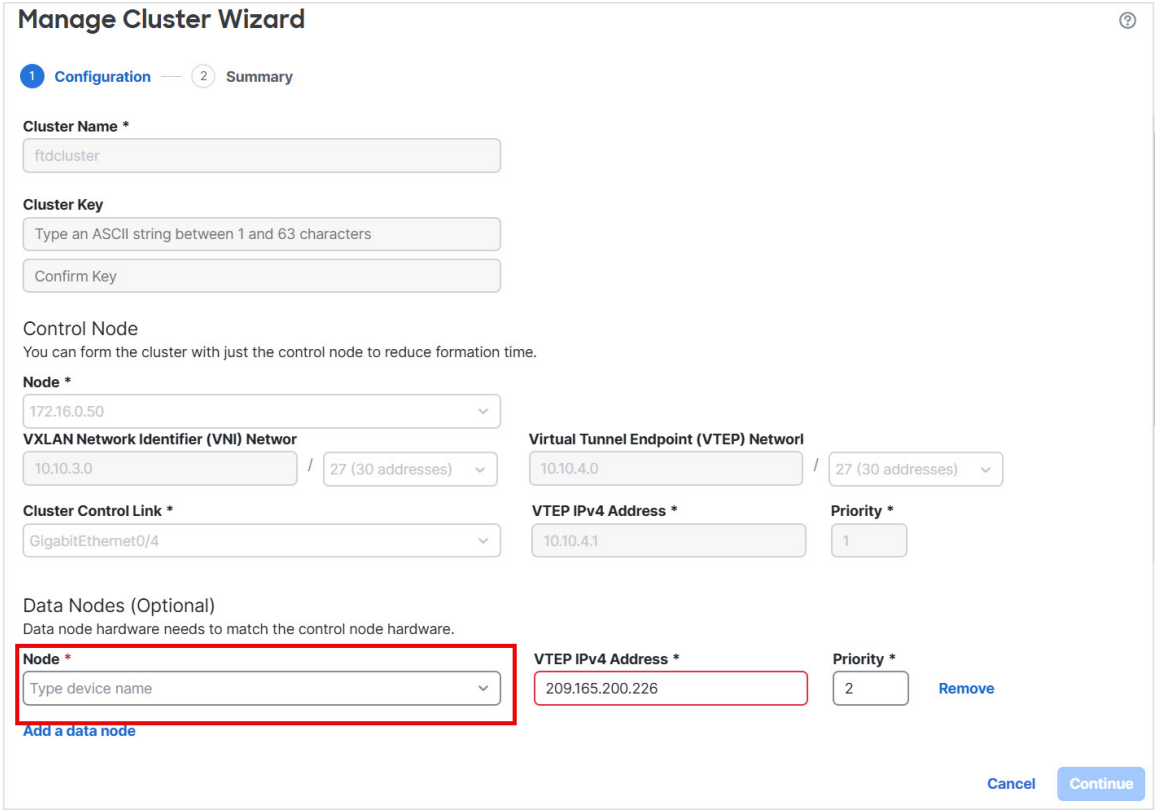

Step 3 To add additional nodes, click **Add a data node**.

Step 4 Click **Continue**. Review the **Summary**, and then click **Save**

The node that is currently registering shows the loading icon.

Figure 248: Node Registration

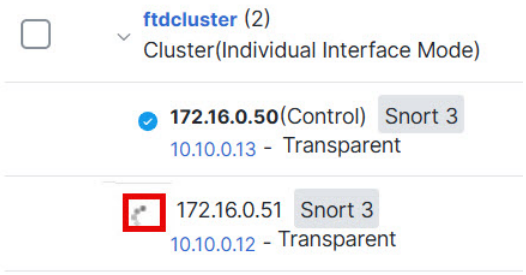

You can monitor cluster node registration by clicking the **Notifications** icon and choosing **Tasks**.

Break a Node

You can remove a node from the cluster so that it becomes a standalone device. You cannot break the control node unless you break the entire cluster. The data node has its configuration erased.

Procedure

Step 1 Choose **Devices** > **Device Management**, click the **More** (⋮) for the node you want to break, and choose **Break Node**.

Figure 249: Break a Node

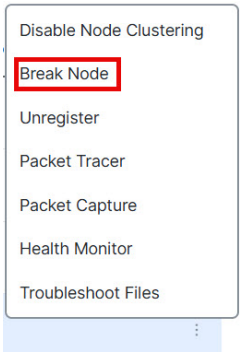

You can optionally break one or more nodes from the cluster More menu by choosing **Break Nodes**.

Step 2 You are prompted to confirm the break; click **Yes**.

Figure 250: Confirm Break

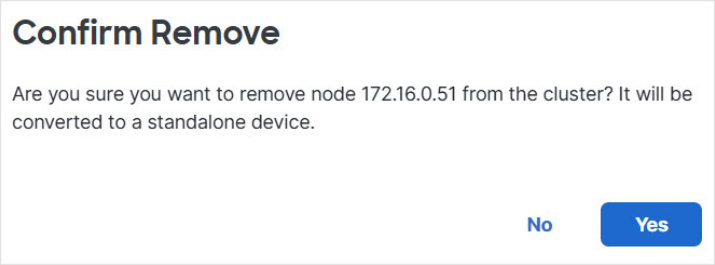

You can monitor the cluster node break by clicking the **Notifications** icon and choosing **Tasks**.

Break the Cluster

You can break the cluster and convert all nodes to standalone devices. The control node retains the interface and security policy configuration, while data nodes have their configuration erased.

Procedure

Step 1 Make sure all cluster nodes are being managed by the Firewall Management Center by reconciling nodes. See Reconcile Cluster Nodes, on page 554.

Step 2 Choose **Devices** > **Device Management**, click the **More** (⋮) for the cluster, and choose **Break Cluster**.

Figure 251: Break Cluster

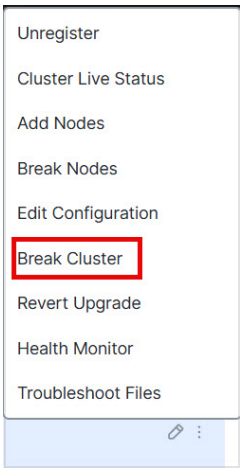

Step 3 You are prompted to break the cluster; click **Yes**.

Figure 252: Confirm Break

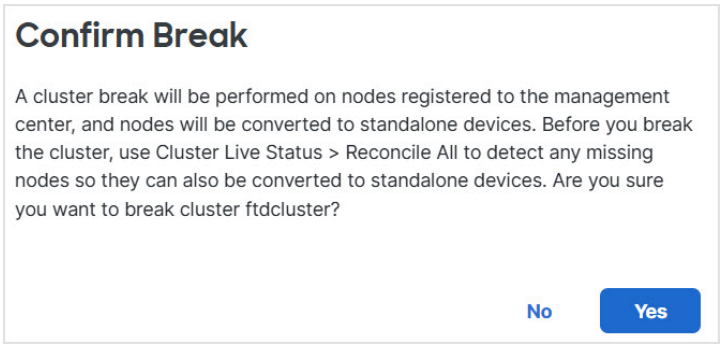

You can monitor the cluster break by clicking the **Notifications** icon and choosing **Tasks**.

Disable Clustering

You may want to deactivate a node in preparation for deleting the node, or temporarily for maintenance. This procedure is meant to temporarily deactivate a node; the node will still appear in the Firewall Management Center device list. When a node becomes inactive, all data interfaces are shut down.

Procedure

Step 1 For the unit you want to disable, choose **Devices** > **Device Management**, click the **More** (⋮), and choose **Disable Node Clustering**.

Figure 253: Disable Clustering

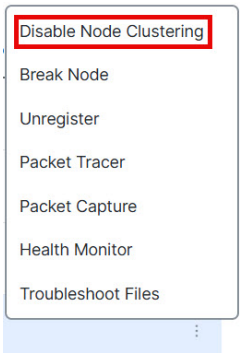

If you disable clustering on the control node, one of the data nodes will become the new control node. Note that for centralized features, if you force a control node change, then all connections are dropped, and you have to re-establish the connections on the new control node. You cannot disable clustering on the control node if it is the only node in the cluster.

Step 2 Confirm that you want to disable clustering on the node.

The node will show **(Disabled)** next to its name in the **Devices** > **Device Management** list.

Step 3 To reenable clustering, see Rejoin the Cluster, on page 551.

Rejoin the Cluster

If a node was removed from the cluster, for example for a failed interface or if you manually disabled clustering, you must manually rejoin the cluster. Make sure the failure is resolved before you try to rejoin the cluster. See Rejoining the Cluster, on page 571 for more information about why a node can be removed from a cluster.

Procedure

Step 1 For the unit you want to reactivate, choose **Devices** > **Device Management**, click the **More** (⋮), and choose **Enable Node Clustering**.

Step 2 Confirm that you want to enable clustering on the node.

Change the Control Node

Caution The best method to change the control node is to disable clustering on the control node, wait for a new control election, and then re-enable clustering. If you must specify the *exact* unit you want to become the control node, use the procedure in this section. Note that for centralized features, if you force a control node change using either method, then all connections are dropped, and you have to re-establish the connections on the new control node.

To change the control node, perform the following steps.

Procedure

Step 1 Open the **Cluster Status** dialog box by choosing **Devices** > **Device Management More** (⋮) > **Cluster Live Status**.

Figure 254: Cluster Status

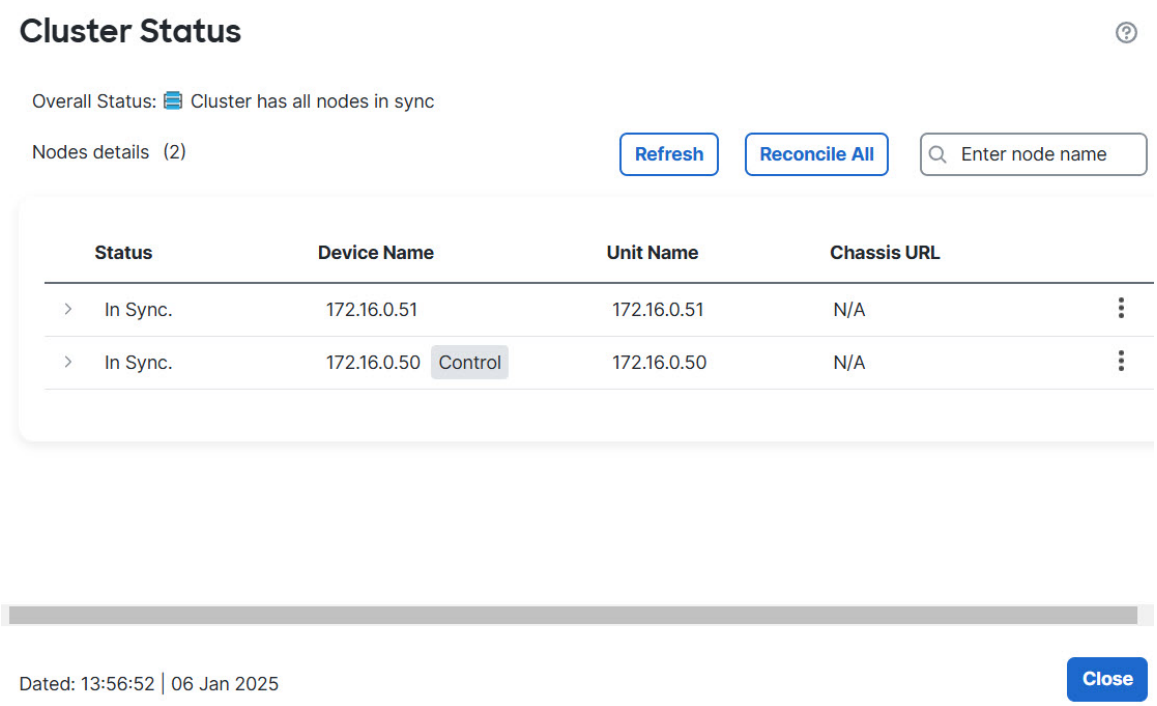

Step 2 For the unit you want to become the control unit, choose **More** (⋮) **> Change Role to Control**.

Step 3 You are prompted to confirm the role change. Check the checkbox, and click **OK**.

Edit the Cluster Configuration

You can edit the cluster configuration. If you change any values other than the VTEP IP address for a node or node priority, the cluster will be broken and reformed automatically. Until the cluster is reformed, you may experience traffic disruption. If you change the VTEP IP address for a node or node priority, only the affected nodes are broken and readded to the cluster.

Procedure

Step 1 Choose **Devices** > **Device Management**, click the **More** (⋮) for the cluster, and choose **Edit Configuration**.

Figure 255: Edit Configuration

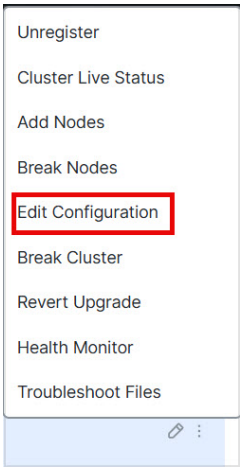

The **Manage Cluster Wizard** appears.

Step 2 Update the cluster configuration.

Figure 256: Manage Cluster Wizard

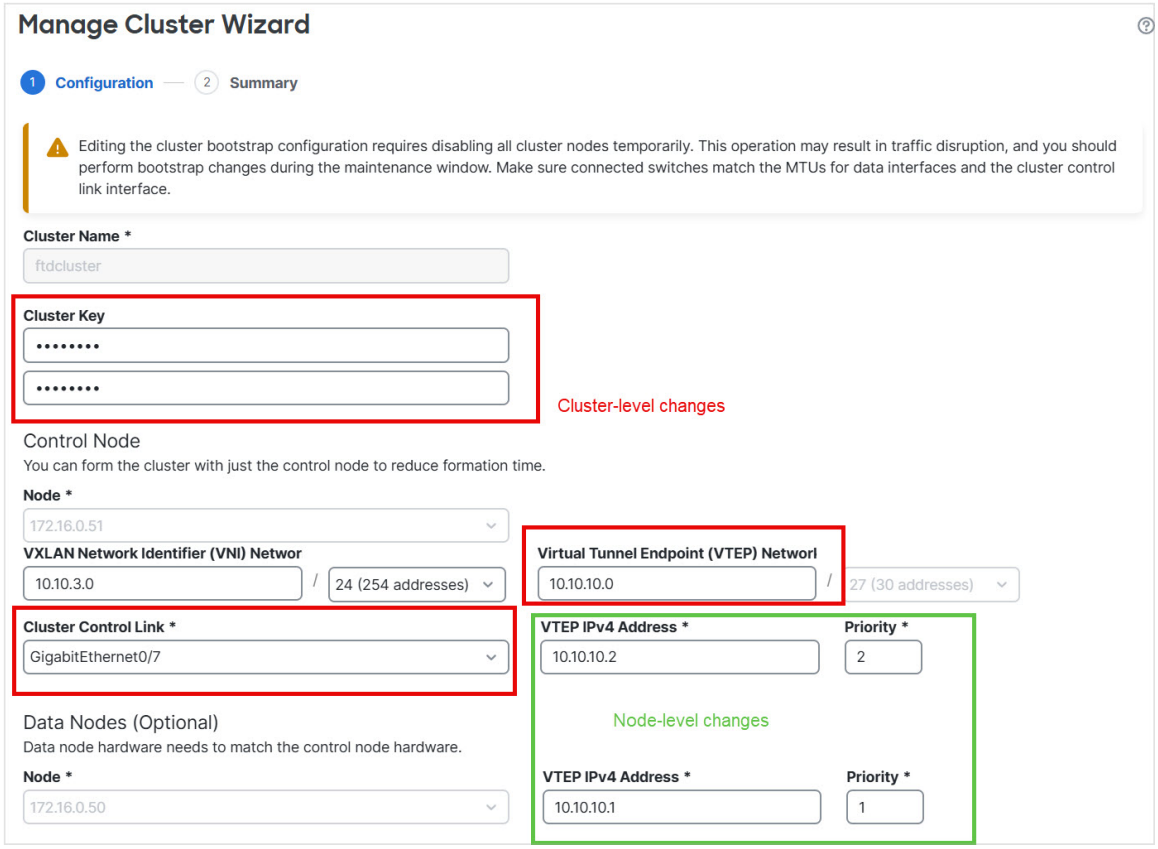

Step 3 Click **Continue**. Review the **Summary**, and then click **Save**

Reconcile Cluster Nodes

If a cluster node fails to register, you can reconcile the cluster membership from the device to the Firewall Management Center. For example, a data node might fail to register if the Firewall Management Center is occupied with certain processes, or if there is a network issue.

Procedure

Step 1 Choose **Devices** > **Device Management More** (⋮) for the cluster, and then choose **Cluster Live Status** to open the **Cluster Status** dialog box.

Figure 257: Cluster Live Status

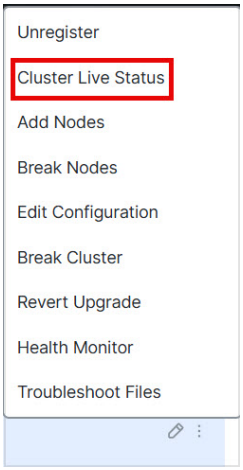

Step 2 Click **Reconcile All**.

Figure 258: Reconcile All

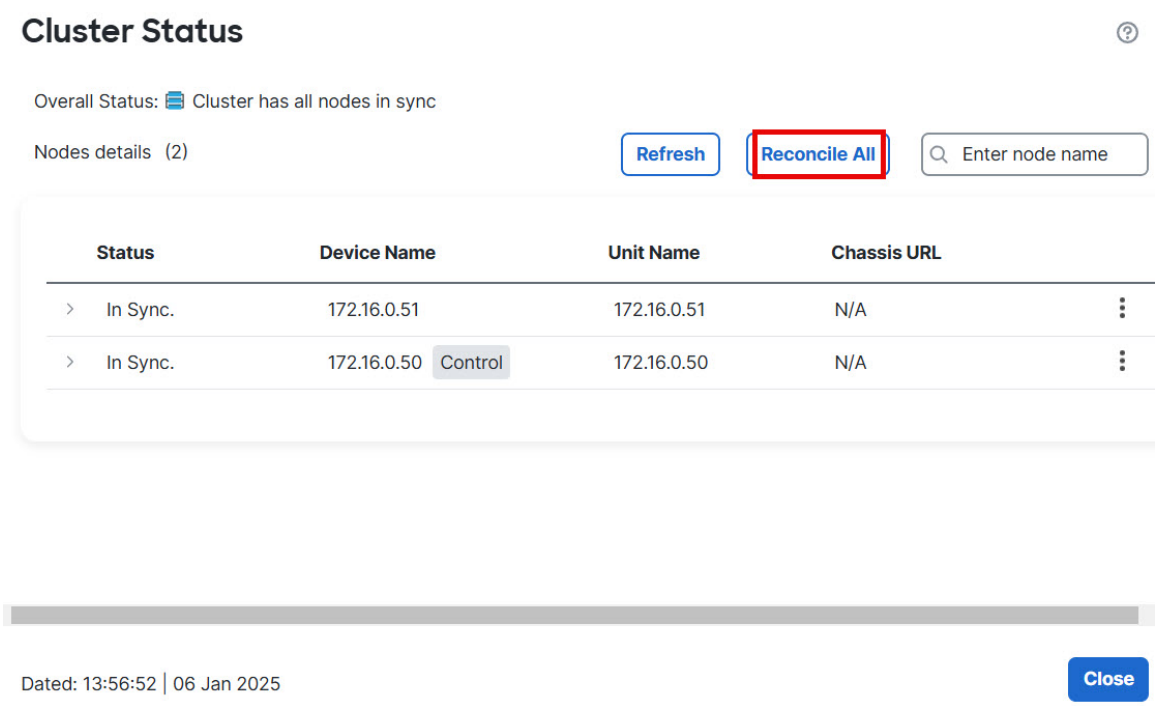

For more information about the cluster status, see Monitoring the Cluster, on page 556.

Delete (Unregister) the Cluster or Nodes and Register to a New Firewall Management Center

You can unregister the cluster from the Firewall Management Center, which keeps the cluster intact. You might want to unregister the cluster if you want to add the cluster to a new Firewall Management Center.

You can also unregister a node from the Firewall Management Center without breaking the node from the cluster. Although the node is not visible in the Firewall Management Center, it is still part of the cluster, and it will continue to pass traffic and could even become the control node. You cannot unregister the current control node. You might want to unregister the node if it is no longer reachable from the Firewall Management Center, but you still want to keep it as part of the cluster while you troubleshoot management connectivity.

Unregistering a cluster:

- Severs all communication between the Firewall Management Center and the cluster.
- Removes the cluster from the **Device Management** page.
- Returns the cluster to local time management if the cluster's platform settings policy is configured to receive time from the Firewall Management Center using NTP.
- Leaves the configuration intact, so the cluster continues to process traffic.

  Policies, such as NAT and VPN, ACLs, and the interface configurations remain intact.

Registering the cluster again to the same or a different Firewall Management Center causes the configuration to be removed, so the cluster will stop processing traffic at that point; the cluster configuration remains intact so you can add the cluster as a whole. You can choose an access control policy at registration, but you will have to re-apply other policies after registration and then deploy the configuration before it will process traffic again.

Before you begin

This procedure requires CLI access to one of the nodes.

Procedure

Step 1 Choose **Devices** > **Device Management**, click **More** (⋮) for the cluster or node, and choose **Unregister**.

Figure 259: Unregister Cluster or Node

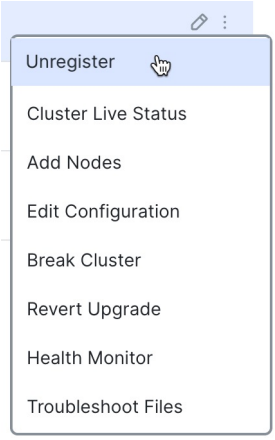

Step 2 You are prompted to unregister the cluster or node; click **Yes**.

Step 3 You can register the cluster to a new (or the same) Firewall Management Center by adding one of the cluster members as a new device.

a) Connect to one cluster node's CLI, and identify the new Firewall Management Center using the **configure manager add** command. See Modify Threat Defense Management Interfaces at the CLI.

b) Choose **Devices** > **Device Management**, and then click **Add Device**.

You only need to add one of the cluster nodes as a device, and the rest of the cluster nodes will be discovered.

Step 4 To re-add a deleted node, see Reconcile Cluster Nodes, on page 554.

Monitoring the Cluster

You can monitor the cluster in the Firewall Management Center and at the Firewall Threat Defense CLI.

- **Cluster Status** dialog box, which is available from the **Devices** > **Device Management More** (⋮) icon or from the **Devices** > **Device Management** > **Cluster** page **> General** area **> Cluster Live Status** link.

Figure 260: Cluster Status

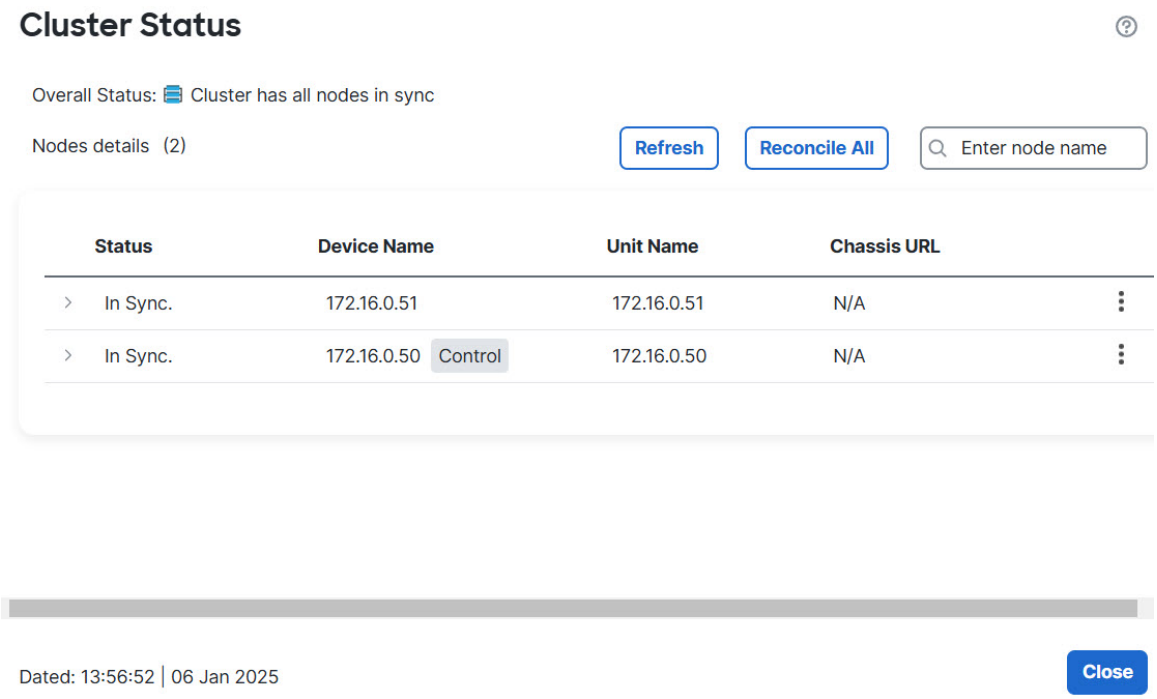

The Control node has a graphic indicator identifying its role.

Cluster member **Status** includes the following states:

- In Sync.—The node is registered with the Firewall Management Center.
- Pending Registration—The node is part of the cluster, but has not yet registered with the Firewall Management Center. If a node fails to register, you can retry registration by clicking **Reconcile All**.
- Clustering is disabled—The node is registered with the Firewall Management Center, but is an inactive member of the cluster. The clustering configuration remains intact if you intend to later re-enable it, or you can delete the node from the cluster.
- Joining cluster...—The node is joining the cluster on the chassis, but has not completed joining. After it joins, it will register with the Firewall Management Center.

For each node, you can view the **Summary** or the **History**.

Figure 261: Node Summary

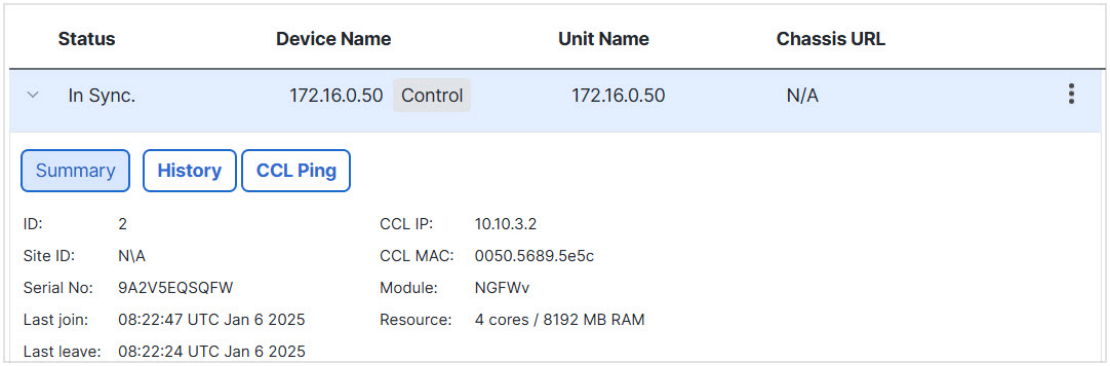

Figure 262: Node History

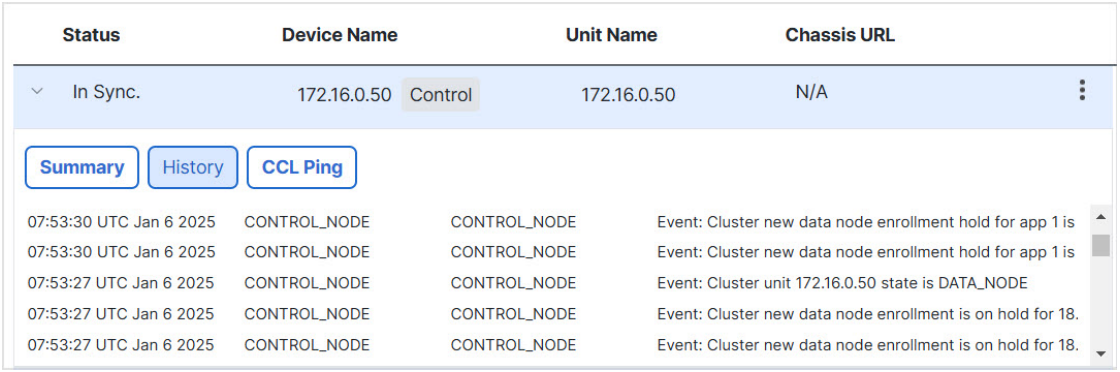

- **System** (⚙) > **Tasks** page.

  The **Tasks** page shows updates of the Cluster Registration task as each node registers.

- **Devices** > **Device Management***cluster_name*.

  When you expand the cluster on the devices listing page, you can see all member nodes, including the control node shown with its role next to the IP address. For nodes that are still registering, you can see the loading icon.

- **show cluster** {**access-list** [*acl_name*] | **conn** [**count**] | **cpu** [**usage**] | **history** | **interface-mode** | **memory** | **resource usage** | **service-policy** | **traffic** | **xlate count**}

  To view aggregated data for the entire cluster or other information, use the **show cluster** command.

- **show cluster info** [**auto-join** | **clients** | **conn-distribution** | **flow-mobility counters** | **goid** [*options*] | **health** | **incompatible-config** | **loadbalance** | **old-members** | **packet-distribution** | **trace** [*options*] | **transport** { **asp** | **cp**}]

  To view cluster information, use the **show cluster info** command.

Cluster Health Monitor Dashboard

Cluster Health Monitor

When a Firewall Threat Defense is the control node of a cluster, the Firewall Management Center collects various metrics periodically from the device metric data collector. The cluster health monitor is comprised of the following components:

- Overview dashboard—Displays information about the cluster topology, cluster statistics, and metric charts:
  - The topology section displays a cluster's live status, the health of individual threat defense, threat defense node type (control node or data node), and the status of the device. The status of the device could be *Disabled* (when the device leaves the cluster), *Added out of box* (in a public cloud cluster, the additional nodes that do not belong to the Firewall Management Center), or *Normal* (ideal state of the node).
  - The cluster statistics section displays current metrics of the cluster with respect to the CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.

Note The CPU and memory metrics display the individual average of the data plane and snort usage.

  - The metric charts, namely, CPU Usage, Memory Usage, Throughput, and Connections, diagrammatically display the statistics of the cluster over the specified time period.

- Load Distribution dashboard—Displays load distribution across the cluster nodes in two widgets:
  - The Distribution widget displays the average packet and connection distribution over the time range across the cluster nodes. This data depicts how the load is being distributed by the nodes. Using this widget, you can easily identify any abnormalities in the load distribution and rectify it.
  - The Node Statistics widget displays the node level metrics in table format. It displays metric data on CPU usage, memory usage, input rate, output rate, active connections, and NAT translations across the cluster nodes. This table view enables you to correlate data and easily identify any discrepancies.
- Member Performance dashboard—Displays current metrics of the cluster nodes. You can use the selector to filter the nodes and view the details of a specific node. The metric data include CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.
- CCL dashboard—Displays, graphically, the cluster control link data namely, the input, and output rate.
- Troubleshooting and Links — Provides convenient links to frequently used troubleshooting topics and procedures.
- Time range—An adjustable time window to constrain the information that appears in the various cluster metrics dashboards and widgets.
- Custom Dashboard—Displays data on both cluster-wide metrics and node-level metrics. However, node selection only applies for the threat defense metrics and not for the entire cluster to which the node belongs.

Viewing Cluster Health

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

The cluster health monitor provides a detailed view of the health status of a cluster and its nodes. This cluster health monitor provides health status and trends of the cluster in an array of dashboards.

Before you begin

- Ensure you have created a cluster from one or more devices in the Firewall Management Center.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Use the Monitoring navigation pane to access node-specific health monitors.

Step 2 In the device list, click **Expand**(>) and **Collapse** (⋁) to expand and collapse the list of managed cluster devices.

Step 3 To view the cluster health statistics, click on the cluster name. The cluster monitor reports health and performance metrics in several predefined dashboards by default. The metrics dashboards include:

- Overview — Highlights key metrics from the other predefined dashboards, including its nodes, CPU, memory, input and output rates, connection statistics, and NAT translation information.
- Load Distribution — Traffic and packet distribution across the cluster nodes.
- Member Performance — Node-level statistics on CPU usage, memory usage, input throughput, output throughput, active connection, and NAT translation.
- CCL — Interface status and aggregate traffic statistics.

You can navigate through the various metrics dashboards by clicking on the labels. For a comprehensive list of the supported cluster metrics, see Cisco Secure Firewall Threat Defense Health Metrics.

Step 4 You can configure the time range from the drop-down in the upper-right corner. The time range can reflect a period as short as the last hour (the default) or as long as two weeks. Select **Custom** from the drop-down to configure a custom start and end date.

Click the refresh icon to set auto refresh to 5 minutes or to toggle off auto refresh.

Step 5 Click on deployment icon for a deployment overlay on the trend graph, with respect to the selected time range.

The deployment icon indicates the number of deployments during the selected time-range. A vertical band indicates the deployment start and end time. For multiple deployments, multiple bands/lines appear. Click on the icon on top of the dotted line to view the deployment details.

Step 6 (For node-specific health monitor) View the **Health Alerts** for the node in the alert notification at the top of page, directly to the right of the device name.

Hover your pointer over the **Health Alerts** to view the health summary of the node. The popup window shows a truncated summary of the top five health alerts. Click on the popup to open a detailed view of the health alert summary.

Step 7 (For node-specific health monitor) The device monitor reports health and performance metrics in several predefined dashboards by default. The metrics dashboards include:

- Overview — Highlights key metrics from the other predefined dashboards, including CPU, memory, interfaces, connection statistics; plus disk usage and critical process information.
- CPU — CPU utilization, including the CPU usage by process and by physical cores.
- Memory — Device memory utilization, including data plane and Snort memory usage.
- Interfaces — Interface status and aggregate traffic statistics.
- Connections — Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
- Snort — Statistics that are related to the Snort process.
- ASP drops — Statistics related to the dropped packets against various reasons.

You can navigate through the various metrics dashboards by clicking on the labels. See Cisco Secure Firewall Threat Defense Health Metrics for a comprehensive list of the supported device metrics.

Step 8 Click the plus sign **Add New Dashboard**(+) in the upper right corner of the health monitor to create a custom dashboard by building your own variable set from the available metric groups.

For cluster-wide dashboard, choose Cluster metric group, and then choose the metric.

Cluster Metrics

The cluster health monitor tracks statistics that are related to a cluster and its nodes, and aggregate of load distribution, performance, and CCL traffic statistics.

Table 35: Cluster Metrics

Metric	Description	Format
CPU	Average of CPU metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Memory	Average of memory metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Data Throughput	Incoming and outgoing data traffic statistics for a cluster.	bytes
CCL Throughput	Incoming and outgoing CCL traffic statistics for a cluster.	bytes
Connections	Count of active connections in a cluster.	number
NAT Translations	Count of NAT translations for a cluster.	number
Distribution	Connection distribution count in the cluster for every second.	number

Metric	Description	Format
Packets	Packet distribution count in the cluster for every second.	number

Troubleshooting the Cluster

You can use the **CCL Ping** tool to make sure the cluster control link is operating correctly. You can also use the following tools that are available for devices and clusters:

- Troubleshooting files—If a node fails to join the cluster, a troubleshooting file is automatically generated. You can also generate and download troubleshooting files from the **Devices** > **Device Management** > **Cluster** > **General** area. See Generate Troubleshooting Files, on page 129.

  You can also generate files from the **Device Management** page by clicking **More** (⋮) and choosing **Troubleshoot Files**.

- CLI output—From the **Devices** > **Device Management** > **Cluster** > **General** area, you can view a set of pre-defined CLI outputs that can help you troubleshoot the cluster. The following commands are automatically run for the cluster:
  - **show running-config cluster**
  - **show cluster info**
  - **show cluster info health**
  - **show cluster info transport cp**
  - **show version**
  - **show asp drop**
  - **show counters**
  - **show arp**
  - **show int ip brief**
  - **show blocks**
  - **show cpu detailed**
  - **show interface** *ccl_interface*
  - **ping** *ccl_ip* **size** *ccl_mtu* **repeat 2**

  You can also enter any **show** command in the Command field. See View CLI Output, on page 132 for more information.

Perform a Ping on the Cluster Control Link

When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the ping fails, a notification is generated so you can fix

the MTU mismatch on connecting switches and try again. This tool lets you manually ping all nodes that have already joined the cluster in case you are having cluster control link connectivity problems.

Procedure

Step 1 Choose **Devices** > **Device Management**, click the **More** (⋮) icon next to the cluster, and choose **Cluster Live Status**.

Figure 263: Cluster Status

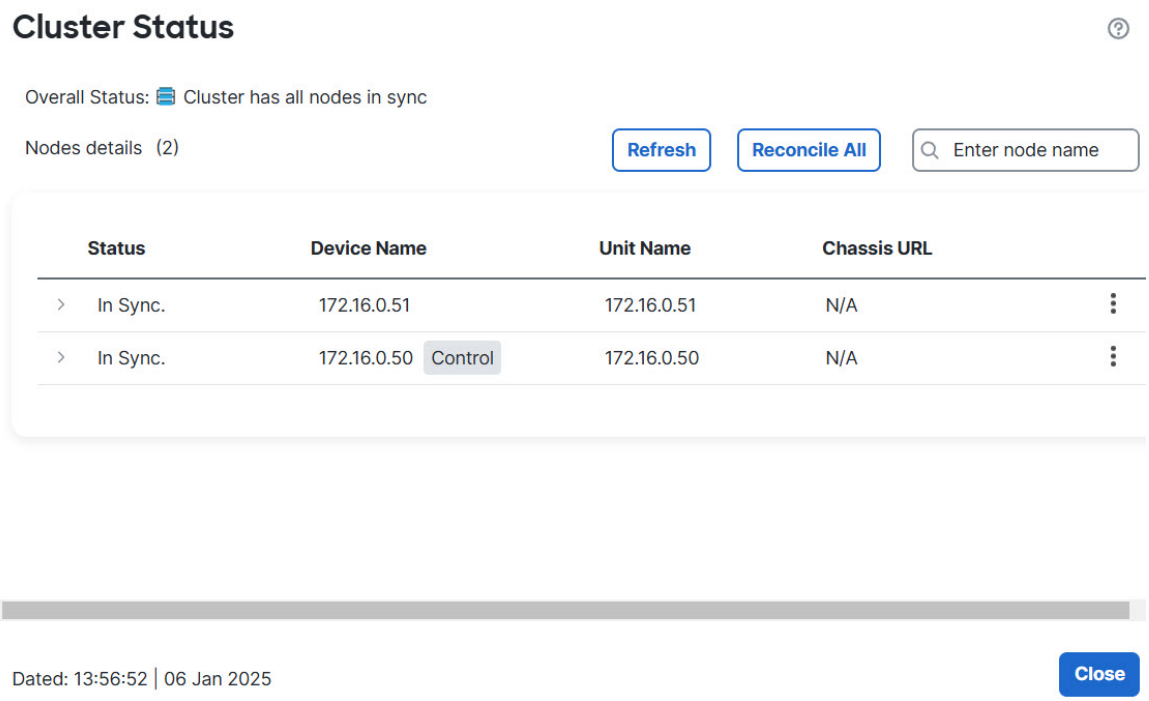

Step 2 Expand one of the nodes, and click **CCL Ping**.

Figure 264: CCL Ping

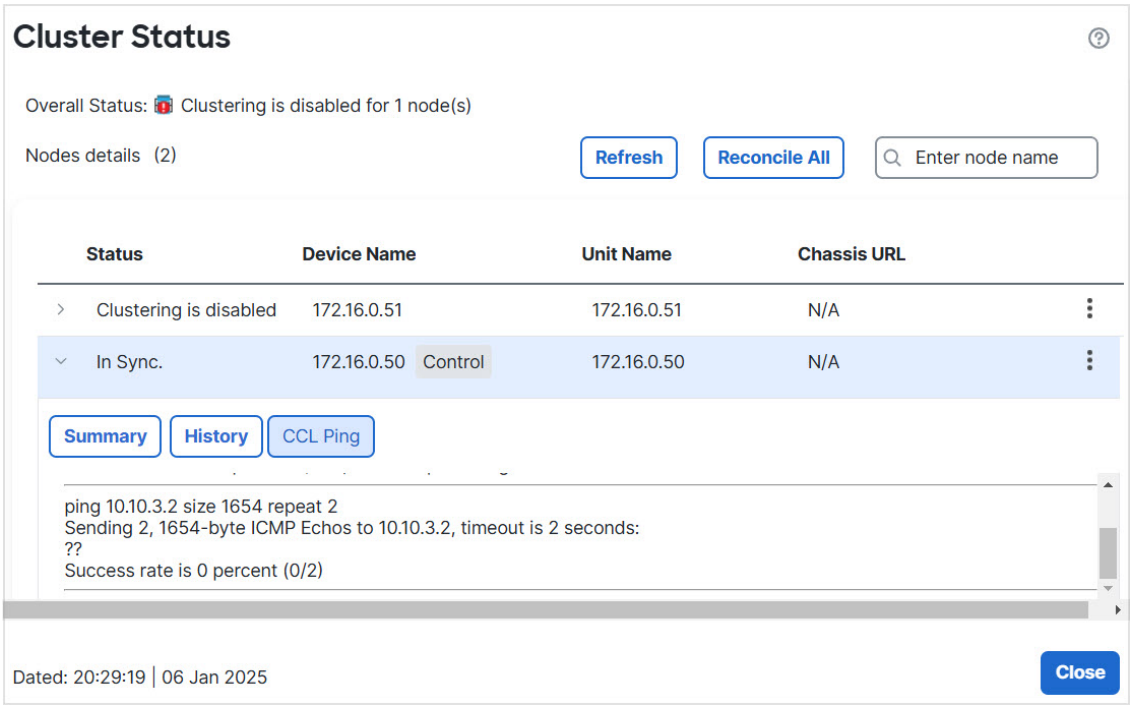

The node sends a ping on the cluster control link to every other node using a packet size that matches the maximum MTU.

Reference for Clustering

This section includes more information about how clustering operates.

Threat Defense Features and Clustering

Some Firewall Threat Defense features are not supported with clustering, and some are only supported on the control unit. Other features might have caveats for proper usage.

Unsupported Features and Clustering

These features cannot be configured with clustering enabled, and the commands will be rejected.

Note To view FlexConfig features that are also not supported with clustering, for example WCCP inspection, see the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the Firewall Management Center GUI. See FlexConfig Policies, on page 2667.

- Remote access VPN (SSL VPN and IPsec VPN)

- DHCP client, server, and proxy. DHCP relay is supported.
- Virtual Tunnel Interfaces (VTIs)
- High Availability
- Integrated Routing and Bridging
- Firewall Management Center UCAPL/CC mode
- DHCP client, server, and proxy. DHCP relay is supported.

Centralized Features for Clustering

The following features are only supported on the control node, and are not scaled for the cluster.

Note Traffic for centralized features is forwarded from member nodes to the control node over the cluster control link.

If you use the rebalancing feature, traffic for centralized features may be rebalanced to non-control nodes before the traffic is classified as a centralized feature; if this occurs, the traffic is then sent back to the control node.

For centralized features, if the control node fails, all connections are dropped, and you have to re-establish the connections on the new control node.

Note To view FlexConfig features that are also centralized with clustering, for example RADIUS inspection, see the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the Firewall Management Center GUI. See FlexConfig Policies, on page 2667.

- The following application inspections:
  - DCERPC
  - ESMTP
  - NetBIOS
  - PPTP
  - RSH
  - SQLNET
  - SUNRPC
  - TFTP
  - XDMCP
- Static route monitoring

Connection Settings and Clustering

Connection limits are enforced cluster-wide. Each node has an estimate of the cluster-wide counter values based on broadcast messages. Due to efficiency considerations, the configured connection limit across the cluster might not be enforced exactly at the limit number. Each node may overestimate or underestimate the cluster-wide counter value at any given time. However, the information will get updated over time in a load-balanced cluster.

Dynamic Routing and Clustering

In Individual interface mode, each node runs the routing protocol as a standalone router, and routes are learned by each node independently.

Figure 265: Dynamic Routing in Individual Interface Mode

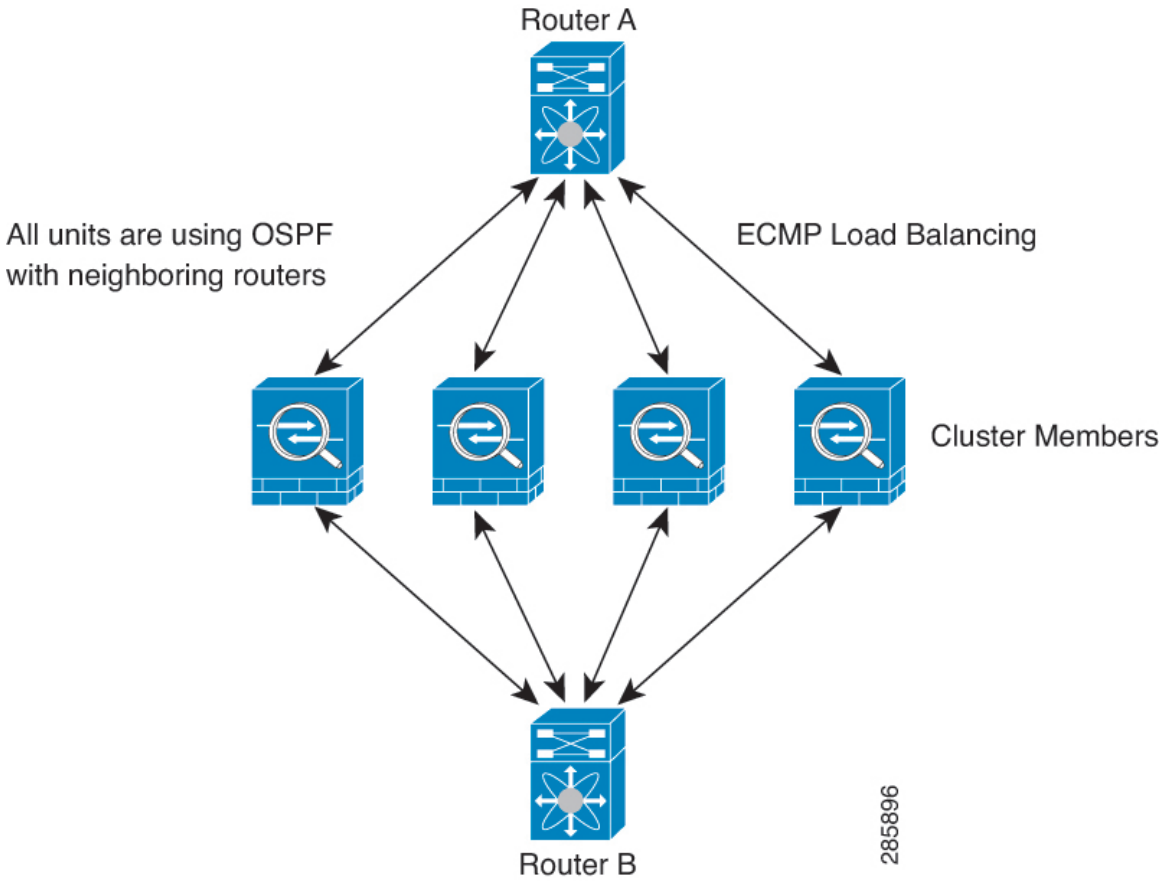

In the above diagram, Router A learns that there are 4 equal-cost paths to Router B, each through a node. ECMP is used to load balance traffic between the 4 paths. Each node picks a different router ID when talking to external routers.

You must configure a cluster pool for the router ID so that each node has a separate router ID.

FTP and Clustering

- If FTP data channel and control channel flows are owned by different cluster members, then the data channel owner will periodically send idle timeout updates to the control channel owner and update the

idle timeout value. However, if the control flow owner is reloaded, and the control flow is re-hosted, the parent/child flow relationship will not longer be maintained; the control flow idle timeout will not be updated.

NAT and Clustering

NAT can affect the overall throughput of the cluster. Inbound and outbound NAT packets can be sent to different Firewall Threat Defenses in the cluster, because the load balancing algorithm relies on IP addresses and ports, and NAT causes inbound and outbound packets to have different IP addresses and/or ports. When a packet arrives at the Firewall Threat Defense that is not the NAT owner, it is forwarded over the cluster control link to the owner, causing large amounts of traffic on the cluster control link. Note that the receiving node does not create a forwarding flow to the owner, because the NAT owner may not end up creating a connection for the packet depending on the results of security and policy checks.

If you still want to use NAT in clustering, then consider the following guidelines:

- No Proxy ARP—For Individual interfaces, a proxy ARP reply is never sent for mapped addresses. This prevents the adjacent router from maintaining a peer relationship with an ASA that may no longer be in the cluster. The upstream router needs a static route or PBR with Object Tracking for the mapped addresses that points to the Main cluster IP address.
- No interface PAT on an Individual interface—Interface PAT is not supported for Individual interfaces.
- PAT with Port Block Allocation—See the following guidelines for this feature:
  - Maximum-per-host limit is not a cluster-wide limit, and is enforced on each node individually. Thus, in a 3-node cluster with the maximum-per-host limit configured as 1, if the traffic from a host is load-balanced across all 3 nodes, then it can get allocated 3 blocks with 1 in each node.
  - Port blocks created on the backup node from the backup pools are not accounted for when enforcing the maximum-per-host limit.
  - On-the-fly PAT rule modifications, where the PAT pool is modified with a completely new range of IP addresses, will result in xlate backup creation failures for the xlate backup requests that were still in transit while the new pool became effective. This behavior is not specific to the port block allocation feature, and is a transient PAT pool issue seen only in cluster deployments where the pool is distributed and traffic is load-balanced across the cluster nodes.
  - When operating in a cluster, you cannot simply change the block allocation size. The new size is effective only after you reload each device in the cluster. To avoid having to reload each device, we recommend that you delete all block allocation rules and clear all xlates related to those rules. You can then change the block size and recreate the block allocation rules.
- NAT pool address distribution for dynamic PAT—When you configure a PAT pool, the cluster divides each IP address in the pool into port blocks. By default, each block is 512 ports, but if you configure port block allocation rules, your block setting is used instead. These blocks are distributed evenly among the nodes in the cluster, so that each node has one or more blocks for each IP address in the PAT pool. Thus, you could have as few as one IP address in a PAT pool for a cluster, if that is sufficient for the number of PAT'ed connections you expect. Port blocks cover the 1024-65535 port range, unless you configure the option to include the reserved ports, 1-1023, on the PAT pool NAT rule.
- Reusing a PAT pool in multiple rules—To use the same PAT pool in multiple rules, you must be careful about the interface selection in the rules. You must either use specific interfaces in all rules, or "any" in all rules. You cannot mix specific interfaces and "any" across the rules, or the system might not be able

to match return traffic to the right node in the cluster. Using unique PAT pools per rule is the most reliable option.

- No round-robin—Round-robin for a PAT pool is not supported with clustering.
- No extended PAT—Extended PAT is not supported with clustering.
- Dynamic NAT xlates managed by the control node—The control node maintains and replicates the xlate table to data nodes. When a data node receives a connection that requires dynamic NAT, and the xlate is not in the table, it requests the xlate from the control node. The data node owns the connection.
- Stale xlates—The xlate idle time on the connection owner does not get updated. Thus, the idle time might exceed the idle timeout. An idle timer value higher than the configured timeout with a refcnt of 0 is an indication of a stale xlate.
- No static PAT for the following inspections—
  - FTP
  - RSH
  - SQLNET
  - TFTP
  - XDMCP
  - SIP
- If you have an extremely large number of NAT rules, over ten thousand, you should enable the transactional commit model using the **asp rule-engine transactional-commit nat** command in the device CLI. Otherwise, the node might not be able to join the cluster.

SIP Inspection and Clustering

A control flow can be created on any node (due to load balancing); its child data flows must reside on the same node.

SNMP and Clustering

You should always use the Local address, and not the Main cluster IP address for SNMP polling. If the SNMP agent polls the Main cluster IP address, if a new control node is elected, the poll to the new control node will fail.

When using SNMPv3 with clustering, if you add a new cluster node after the initial cluster formation, then SNMPv3 users are not replicated to the new node. You must remove the users, and re-add them, and then redeploy your configuration to force the users to replicate to the new node.

Syslog and Clustering

- Each node in the cluster generates its own syslog messages. You can configure logging so that each node uses either the same or a different device ID in the syslog message header field. For example, the hostname configuration is replicated and shared by all nodes in the cluster. If you configure logging to use the hostname as the device ID, syslog messages generated by all nodes look as if they come from a single node. If you configure logging to use the local-node name that is assigned in the cluster bootstrap configuration as the device ID, syslog messages look as if they come from different nodes.

Cisco Trustsec and Clustering

Only the control node learns security group tag (SGT) information. The control node then populates the SGT to data nodes, and data nodes can make a match decision for SGT based on the security policy.

VPN and Clustering

VPN functionality is limited to the control node and does not take advantage of the cluster high availability capabilities. If the control node fails, all existing VPN connections are lost, and VPN users will see a disruption in service. When a new control node is elected, you must reestablish the VPN connections.

For connections to an Individual interface when using PBR or ECMP, you must always connect to the Main cluster IP address, not a Local address.

VPN-related keys and certificates are replicated to all nodes.

Note Remote access VPN is not supported with clustering.

Performance Scaling Factor

When you combine multiple units into a cluster, you can expect the total cluster performance to be approximately 80% of the maximum combined throughput.

For example, if your model can handle approximately 10 Gbps of traffic when running alone, then for a cluster of 8 units, the maximum combined throughput will be approximately 80% of 80 Gbps (8 units x 10 Gbps): 64 Gbps.

Control Node Election

Nodes of the cluster communicate over the cluster control link to elect a control node as follows:

1. When you enable clustering for a node (or when it first starts up with clustering already enabled), it broadcasts an election request every 3 seconds.
2. Any other nodes with a higher priority respond to the election request; the priority is set between 1 and 100, where 1 is the highest priority.
3. If after 45 seconds, a node does not receive a response from another node with a higher priority, then it becomes the control node.

Note If multiple nodes tie for the highest priority, the cluster node name and then the serial number is used to determine the control node.

4. If a node later joins the cluster with a higher priority, it does not automatically become the control node; the existing control node always remains as the control node unless it stops responding, at which point a new control node is elected.
5. In a "split brain" scenario when there are temporarily multiple control nodes, then the node with highest priority retains the role while the other nodes return to data node roles.

Note You can manually force a node to become the control node. For centralized features, if you force a control node change, then all connections are dropped, and you have to re-establish the connections on the new control node.

High Availability within the Cluster

Clustering provides high availability by monitoring node and interface health and by replicating connection states between nodes.

Node Health Monitoring

Each node periodically sends a broadcast heartbeat packet over the cluster control link. If the control node does not receive any heartbeat packets or other packets from a data node within the configurable timeout period, then the control node removes the data node from the cluster. If the data nodes do not receive packets from the control node, then a new control node is elected from the remaining nodes.

If nodes cannot reach each other over the cluster control link because of a network failure and not because a node has actually failed, then the cluster may go into a "split brain" scenario where isolated data nodes will elect their own control nodes. For example, if a router fails between two cluster locations, then the original control node at location 1 will remove the location 2 data nodes from the cluster. Meanwhile, the nodes at location 2 will elect their own control node and form their own cluster. Note that asymmetric traffic may fail in this scenario. After the cluster control link is restored, then the control node that has the higher priority will keep the control node’s role.

Interface Monitoring

Each node monitors the link status of all named hardware interfaces in use, and reports status changes to the control node.

All physical interfaces are monitored; only named interfaces can be monitored. You can optionally disable monitoring per interface.

A node is removed from the cluster if its monitored interfaces fail. The node is removed after 500 ms.

Status After Failure

When a node in the cluster fails, the connections hosted by that node are seamlessly transferred to other nodes; state information for traffic flows is shared over the control node's cluster control link.

If the control node fails, then another member of the cluster with the highest priority (lowest number) becomes the control node.

The Firewall Threat Defense automatically tries to rejoin the cluster, depending on the failure event.

Note When the Firewall Threat Defense becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the Management interface can send and receive traffic.

Rejoining the Cluster

After a cluster member is removed from the cluster, how it can rejoin the cluster depends on why it was removed:

- Failed cluster control link when initially joining—After you resolve the problem with the cluster control link, you must manually rejoin the cluster by re-enabling clustering.
- Failed cluster control link after joining the cluster—The Firewall Threat Defense automatically tries to rejoin every 5 minutes, indefinitely.
- Failed data interface—The Firewall Threat Defense automatically tries to rejoin at 5 minutes, then at 10 minutes, and finally at 20 minutes. If the join is not successful after 20 minutes, then the Firewall Threat Defense application disables clustering. After you resolve the problem with the data interface, you have to manually enable clustering.
- Failed node—If the node was removed from the cluster because of a node health check failure, then rejoining the cluster depends on the source of the failure. For example, a temporary power failure means the node will rejoin the cluster when it starts up again as long as the cluster control link is up. The Firewall Threat Defense application attempts to rejoin the cluster every 5 seconds.
- Internal error—Internal failures include: application sync timeout; inconsistent application statuses; and so on.
- Failed configuration deployment—If you deploy a new configuration from Firewall Management Center, and the deployment fails on some cluster members but succeeds on others, then the nodes that failed are removed from the cluster. You must manually rejoin the cluster by re-enabling clustering. If the deployment fails on the control node, then the deployment is rolled back, and no members are removed. If the deployment fails on all data nodes, then the deployment is rolled back, and no members are removed.

Data Path Connection State Replication

Every connection has one owner and at least one backup owner in the cluster. The backup owner does not take over the connection in the event of a failure; instead, it stores TCP/UDP state information, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner is usually also the director.

Some traffic requires state information above the TCP or UDP layer. See the following table for clustering support or lack of support for this kind of traffic.

Table 36: Features Replicated Across the Cluster

Traffic	State Support	Notes
Up time	Yes	Keeps track of the system up time.
ARP Table	Yes	—
MAC address table	Yes	—
User Identity	Yes	—
IPv6 Neighbor database	Yes	—
Dynamic routing	Yes	—

Traffic	State Support	Notes
SNMP Engine ID	**No**	—

How the Cluster Manages Connections

Connections can be load-balanced to multiple nodes of the cluster. Connection roles determine how connections are handled in both normal operation and in a high availability situation.

Connection Roles

See the following roles defined for each connection:

- Owner—Usually, the node that initially receives the connection. The owner maintains the TCP state and processes packets. A connection has only one owner. If the original owner fails, then when new nodes receive packets from the connection, the director chooses a new owner from those nodes.
- Backup owner—The node that stores TCP/UDP state information received from the owner, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner does not take over the connection in the event of a failure. If the owner becomes unavailable, then the first node to receive packets from the connection (based on load balancing) contacts the backup owner for the relevant state information so it can become the new owner.

  As long as the director (see below) is not the same node as the owner, then the director is also the backup owner. If the owner chooses itself as the director, then a separate backup owner is chosen.

  For clustering on the Firepower 9300, which can include up to 3 cluster nodes in one chassis, if the backup owner is on the same chassis as the owner, then an additional backup owner will be chosen from another chassis to protect flows from a chassis failure.
- Director—The node that handles owner lookup requests from forwarders. When the owner receives a new connection, it chooses a director based on a hash of the source/destination IP address and ports (see below for ICMP hash details), and sends a message to the director to register the new connection. If packets arrive at any node other than the owner, the node queries the director about which node is the owner so it can forward the packets. A connection has only one director. If a director fails, the owner chooses a new director.

  As long as the director is not the same node as the owner, then the director is also the backup owner (see above). If the owner chooses itself as the director, then a separate backup owner is chosen.

  ICMP/ICMPv6 hash details:

  - For Echo packets, the source port is the ICMP identifier, and the destination port is 0.
  - For Reply packets, the source port is 0, and the destination port is the ICMP identifier.
  - For other packets, both source and destination ports are 0.
- Forwarder—A node that forwards packets to the owner. If a forwarder receives a packet for a connection it does not own, it queries the director for the owner, and then establishes a flow to the owner for any other packets it receives for this connection. The director can also be a forwarder. Note that if a forwarder receives the SYN-ACK packet, it can derive the owner directly from a SYN cookie in the packet, so it does not need to query the director. (If you disable TCP sequence randomization, the SYN cookie is not used; a query to the director is required.) For short-lived flows such as DNS and ICMP, instead of querying, the forwarder immediately sends the packet to the director, which then sends them to the owner.

A connection can have multiple forwarders; the most efficient throughput is achieved by a good load-balancing method where there are no forwarders and all packets of a connection are received by the owner.

Note We do not recommend disabling TCP sequence randomization when using clustering. There is a small chance that some TCP sessions won't be established, because the SYN/ACK packet might be dropped.

- Fragment Owner—For fragmented packets, cluster nodes that receive a fragment determine a fragment owner using a hash of the fragment source IP address, destination IP address, and the packet ID. All fragments are then forwarded to the fragment owner over the cluster control link. Fragments may be load-balanced to different cluster nodes, because only the first fragment includes the 5-tuple used in the switch load balance hash. Other fragments do not contain the source and destination ports and may be load-balanced to other cluster nodes. The fragment owner temporarily reassembles the packet so it can determine the director based on a hash of the source/destination IP address and ports. If it is a new connection, the fragment owner will register to be the connection owner. If it is an existing connection, the fragment owner forwards all fragments to the provided connection owner over the cluster control link. The connection owner will then reassemble all fragments.

Port Address Translation Connections

New Connection Ownership

Traffic redirection is not supported in this release. When a new connection is directed to a node of the cluster via load balancing, that node owns both directions of the connection. All the subsequent packets for the same connection should arrive the same node. If any connection packets arrive at a different node, they will be dropped. If a reverse flow arrives at a different node, it will be dropped as well. For centralized features, if the connections do not arrive on the control node, they will be dropped.

By default, AWS GWLB uses 5-tuple to maintain flow stickiness. It is recommended to enable 2-tuple or 3-tuple stickiness on AWS GWLB to ensure the same flows are sent to the same node.

Sample Data Flow for TCP

The following example shows the establishment of a new connection.

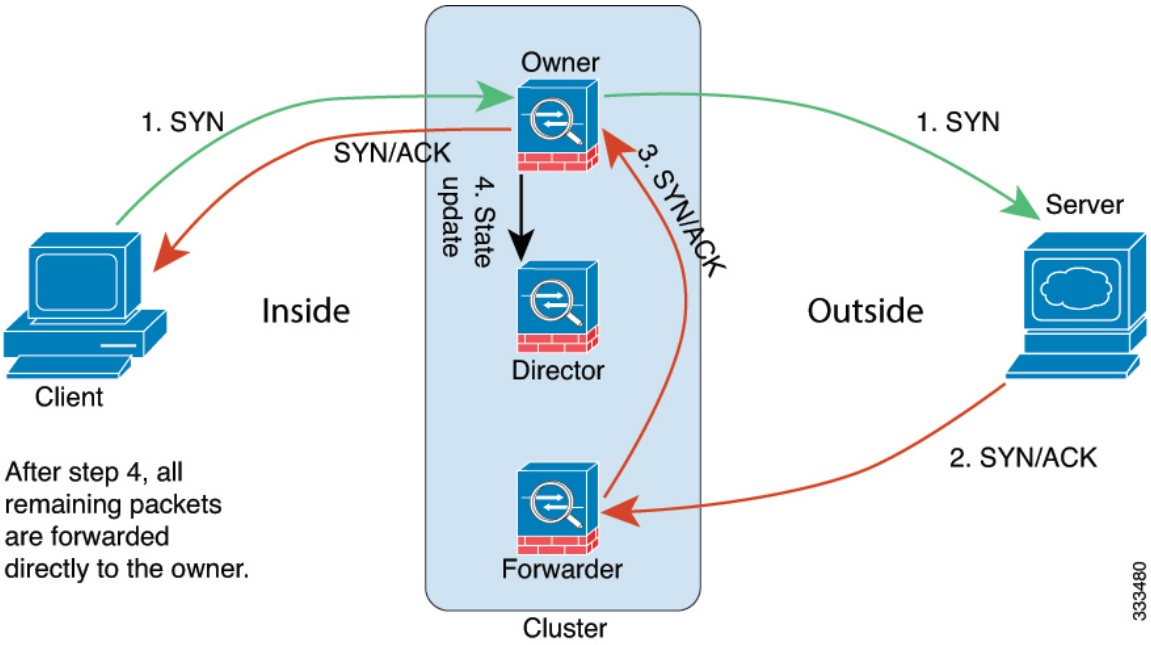

1. The SYN packet originates from the client and is delivered to one Firewall Threat Defense (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.
2. The SYN-ACK packet originates from the server and is delivered to a different Firewall Threat Defense (based on the load balancing method). This Firewall Threat Defense is the forwarder.
3. Because the forwarder does not own the connection, it decodes owner information from the SYN cookie, creates a forwarding flow to the owner, and forwards the SYN-ACK to the owner.
4. The owner sends a state update to the director, and forwards the SYN-ACK to the client.
5. The director receives the state update from the owner, creates a flow to the owner, and records the TCP state information as well as the owner. The director acts as the backup owner for the connection.
6. Any subsequent packets delivered to the forwarder will be forwarded to the owner.
7. If packets are delivered to any additional nodes, it will query the director for the owner and establish a flow.
8. Any state change for the flow results in a state update from the owner to the director.

Sample Data Flow for ICMP and UDP

The following example shows the establishment of a new connection.

1. *Figure 266: ICMP and UDP Data Flow*

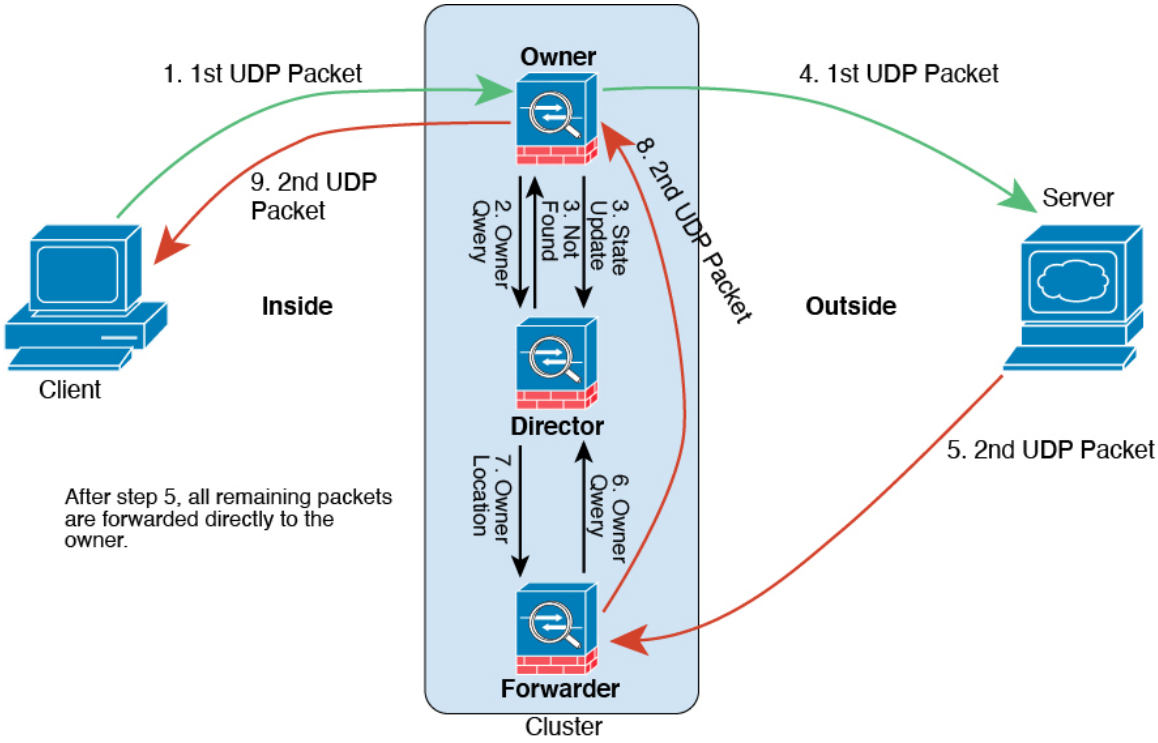

The first UDP packet originates from the client and is delivered to one Firewall Threat Defense (based on the load balancing method).

2. The node that received the first packet queries the director node that is chosen based on a hash of the source/destination IP address and ports.
3. The director finds no existing flow, creates a director flow and forwards the packet back to the previous node. In other words, the director has elected an owner for this flow.
4. The owner creates the flow, sends a state update to the director, and forwards the packet to the server.
5. The second UDP packet originates from the server and is delivered to the forwarder.
6. The forwarder queries the director for ownership information. For short-lived flows such as DNS, instead of querying, the forwarder immediately sends the packet to the director, which then sends it to the owner.
7. The director replies to the forwarder with ownership information.
8. The forwarder creates a forwarding flow to record owner information and forwards the packet to the owner.
9. The owner forwards the packet to the client.

History for Threat Defense Virtual Clustering in a Private Cloud

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
MTU ping test on cluster node join	7.6.0	7.6.0	When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the ping fails, a notification is generated so you can fix the MTU mismatch on connecting switches and try again.
Clustering for the Firewall Threat Defense Virtual on VMware and KVM	7.4.1	7.4.1	The Firewall Threat Defense Virtual now supports individual interface clustering for up to 16 nodes on VMware and KVM.
Cluster control link ping tool.	7.2.6/7.4.1	Any	You can check to make sure all the cluster nodes can reach each other over the cluster control link by performing a ping. One major cause for the failure of a node to join the cluster is an incorrect cluster control link configuration; for example, the cluster control link MTU may be set higher than the connecting switch MTUs. New/modified screens: **Devices** > **Device Management** > **More** > **Cluster Live Status.** Version restrictions: Not supported with Firewall Management Center Version 7.3.x or 7.4.0.
Troubleshooting file generation and download available from Device and Cluster pages.	7.4.1	7.4.1	You can generate and download troubleshooting files for each device on the Device page and also for all cluster nodes on the Cluster page. For a cluster, you can download all files as a single compressed file. You can also include cluster logs for the cluster for cluster nodes. You can alternatively trigger file generation from the **Devices** > **Device Management** > **More** > **Troubleshoot Files** menu. New/modified screens: • **Devices** > **Device Management** > **Device** > **General** • **Devices** > **Device Management** > **Cluster** > **General**
Automatic generation of a troubleshooting file on a node when it fails to join the cluster.	7.4.1	7.4.1	If a node fails to join the cluster, a troubleshooting file is automatically generated for the node. You can download the file from **Tasks** or from the **Cluster** page.
View CLI output for a device or device cluster.	7.4.1	Any	You can view a set of pre-defined CLI outputs that can help you troubleshoot the device or cluster. You can also enter any **show** command and see the output. New/modified screens: **Devices** > **Device Management** > **Cluster** > **General**

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Cluster health monitor settings	7.3.0	Any	You can now edit cluster health monitor settings. New/Modified screens: **Devices** > **Device Management** > **Cluster** > **Cluster Health Monitor Settings** **Note** If you previously configured these settings using FlexConfig, be sure to remove the FlexConfig configuration before you deploy. Otherwise the FlexConfig configuration will overwrite the management center configuration.
Cluster health monitor dashboard	7.3.0	Any	You can now view cluster health on the cluster health monitor dashboard. New/Modified screens: **System** > **Health** > **Monitor**
Clustering for the Firewall Threat Defense Virtual on VMware and KVM	7.2.0	7.2.0	The Firewall Threat Defense Virtual supports Individual interface clustering for up to 4 nodes on VMware and KVM. New/Modified screens: • **Devices** > **Device Management** > **Add Cluster** • **Devices** > **Device Management** > **More** menu • **Devices** > **Device Management** > **Cluster** Supported platforms: Firewall Threat Defense Virtual on VMware and KVM

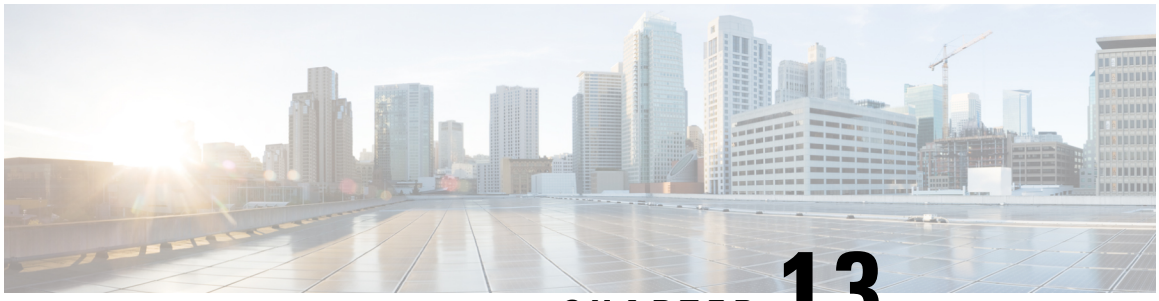

CHAPTER 13

Clustering for Threat Defense Virtual in a Public Cloud

Clustering lets you group multiple Firewall Threat Defense Virtuals together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices. You can deploy Firewall Threat Defense Virtual clusters in a public cloud using the following public cloud platforms:

- Amazon Web Services (AWS)
- Microsoft Azure
- Google Cloud Platform (GCP)

Currently, only routed firewall mode is supported.

Note Some features are not supported when using clustering. See Unsupported Features and Clustering, on page 685.

- About Threat Defense Virtual Clustering in the Public Cloud, on page 580
- Licenses for Threat Defense Virtual Clustering, on page 582
- Requirements and Prerequisites for Threat Defense Virtual Clustering, on page 583
- Guidelines for Threat Defense Virtual Clustering, on page 584
- Deploy the Cluster in AWS, on page 586
- Deploy the Cluster in Azure, on page 610
- Firewall Threat Defense Virtual Clustering Autoscale Solution on Azure, on page 629
- Deploy the Cluster in GCP, on page 653
- Add the Cluster to the Management Center (Manual Deployment), on page 662
- Configure Cluster Health Monitor Settings, on page 669
- Manage Cluster Nodes, on page 674
- Monitoring the Cluster, on page 676
- Troubleshooting the Cluster, on page 682
- Upgrading the Cluster, on page 684
- Reference for Clustering, on page 685
- History for Threat Defense Virtual Clustering in the Public Cloud, on page 696

About Threat Defense Virtual Clustering in the Public Cloud

This section describes the clustering architecture and how it works.

How the Cluster Fits into Your Network

The cluster consists of multiple firewalls acting as a single device. To act as a cluster, the firewalls need the following infrastructure:

- Isolated network for intra-cluster communication, known as the *cluster control link*, using VXLAN interfaces. VXLANs, which act as Layer 2 virtual networks over Layer 3 physical networks, let the Firewall Threat Defense Virtual send broadcast/multicast messages over the cluster control link.
- Load Balancer(s)—For external load balancing, you have the following options depending on your public cloud:
  - AWS Gateway Load Balancer

    The AWS Gateway Load Balancer combines a transparent network gateway and a load balancer that distributes traffic and scales virtual appliances on demand. The Firewall Threat Defense Virtual supports the Gateway Load Balancer centralized control plane with a distributed data plane (Gateway Load Balancer endpoint) using a Geneve interface single-arm proxy.
  - Azure Gateway Load Balancer

    In an Azure service chain, Firewall Threat Defense Virtuals act as a transparent gateway that can intercept packets between the internet and the customer service. The Firewall Threat Defense Virtual defines an external interface and an internal interface on a single NIC by utilizing VXLAN segments in a paired proxy.
  - Native GCP load balancers, internal and external
  - Equal-Cost Multi-Path Routing (ECMP) using inside and outside routers such as Cisco Cloud Services Router

    ECMP routing can forward packets over multiple “best paths” that tie for top place in the routing metric. Like EtherChannel, a hash of source and destination IP addresses and/or source and destination ports can be used to send a packet to one of the next hops. If you use static routes for ECMP routing, then the Firewall Threat Defense failure can cause problems; the route continues to be used, and traffic to the failed Firewall Threat Defense will be lost. If you use static routes, be sure to use a static route monitoring feature such as Object Tracking. We recommend using dynamic routing protocols to add and remove routes, in which case, you must configure each Firewall Threat Defense to participate in dynamic routing.

Note Layer 2 Spanned EtherChannels are not supported for load balancing.

Individual Interfaces

You can configure cluster interfaces as *Individual interfaces*.

Individual interfaces are normal routed interfaces, each with their own local IP address. The IP address for the interface will be configured automatically via DHCP. Static IP configuration is not supported.

Control and Data Node Roles

One member of the cluster is the control node. If multiple cluster nodes come online at the same time, the control node is determined by the priority setting; the priority is set between 1 and 100, where 1 is the highest priority. All other members are data nodes. When you first create the cluster, you specify which node you want to be the control node, and it will become the control node simply because it is the first node added to the cluster.

All nodes in the cluster share the same configuration. The node that you initially specify as the control node will overwrite the configuration on the data nodes when they join the cluster, so you only need to perform initial configuration on the control node before you form the cluster.

Some features do not scale in a cluster, and the control node handles all traffic for those features.

Cluster Control Link

Each node must dedicate one interface as a VXLAN (VTEP) interface for the cluster control link. For more information about VXLAN, see Configure VXLAN Interfaces, on page 848.

VXLAN Tunnel Endpoint

VXLAN tunnel endpoint (VTEP) devices perform VXLAN encapsulation and decapsulation. Each VTEP has two interface types: one or more virtual interfaces called VXLAN Network Identifier (VNI) interfaces, and a regular interface called the VTEP source interface that tunnels the VNI interfaces between VTEPs. The VTEP source interface is attached to the transport IP network for VTEP-to-VTEP communication.

VTEP Source Interface

The VTEP source interface is a regular Firewall Threat Defense Virtual interface with which you plan to associate the VNI interface. You can configure one VTEP source interface to act as the cluster control link. The source interface is reserved for cluster control link use only. Each VTEP source interface has an IP address on the same subnet. This subnet should be isolated from all other traffic, and should include only the cluster control link interfaces.

VNI Interface

A VNI interface is similar to a VLAN interface: it is a virtual interface that keeps network traffic separated on a given physical interface by using tagging. You can only configure one VNI interface. Each VNI interface has an IP address on the same subnet.

Peer VTEPs

Unlike regular VXLAN for data interfaces, which allows a single VTEP peer, The Firewall Threat Defense Virtual clustering allows you to configure multiple peers.

Cluster Control Link Traffic Overview

Cluster control link traffic includes both control and data traffic.

Control traffic includes:

- Control node election.
- Configuration replication.
- Health monitoring.

Data traffic includes:

- State replication.
- Connection ownership queries and data packet forwarding.

Configuration Replication

All nodes in the cluster share a single configuration. You can only make configuration changes on the control node (with the exception of the bootstrap configuration), and changes are automatically synced to all other nodes in the cluster.

Management Network

You must manage each node using the Management interface; management from a data interface is not supported with clustering.

Licenses for Threat Defense Virtual Clustering

Each Firewall Threat Defense Virtual cluster node requires the same performance tier license. We recommend using the same number of CPUs and memory for all members, or else performance will be limited on all nodes to match the least capable member. The throughput level will be replicated from the control node to each data node so they match.

You assign feature licenses to the cluster as a whole, not to individual nodes. However, each node of the cluster consumes a separate license for each feature. The clustering feature itself does not require any licenses.

When you add the control node to the Firewall Management Center, you can specify the feature licenses you want to use for the cluster. You can modify licenses for the cluster in the **Devices** > **Device Management** > **Cluster** > **License** area.

Note If you add the cluster before the Firewall Management Center is licensed (and running in Evaluation mode), then when you license the Firewall Management Center, you can experience traffic disruption when you deploy policy changes to the cluster. Changing to licensed mode causes all data units to leave the cluster and then rejoin.

Requirements and Prerequisites for Threat Defense Virtual Clustering

Model Requirements

- FTDv5, FTDv10, FTDv20, FTDv30, FTDv50, FTDv100

Note FTDv5 and FTDv10 do not support Amazon Web Services (AWS) Gateway Load Balancer (GWLB) and Azure GWLB.

- The following public cloud services:
  - Amazon Web Services (AWS)
  - Microsoft Azure
  - Google Cloud Platform (GCP)
- Maximum 16 nodes

See also the general requirements for the Firewall Threat Defense Virtual in the Cisco Secure Firewall Threat Defense Virtual Getting Started Guide.

User Roles

- Admin
- Access Admin
- Network Admin

Hardware and Software Requirements

All units in a cluster:

- Must be in the same performance tier. We recommend using the same number of CPUs and memory for all nodes, or else performance will be limited on all nodes to match the least capable node.
- The Firewall Management Center access must be from the Management interface; data interface management is not supported.
- Must run the identical software except at the time of an image upgrade. Hitless upgrade is supported.
- All units in a cluster must be deployed in the same availability zone.
- Cluster control link interfaces of all units must be in the same subnet.

MTU

Make sure the ports connected to the cluster control link have the correct (higher) MTU configured. If there is an MTU mismatch, the cluster formation will fail. When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the ping fails, a notification is generated so you can fix the MTU mismatch on connecting switches and try again.

The cluster control link MTU should be 154 bytes higher than the data interfaces. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead (100 bytes) plus VXLAN overhead (54 bytes).

For AWS with GWLB, the data interface uses Geneve encapsulation. In this case, the entire Ethernet datagram is being encapsulated, so the new packet is larger and requires a larger MTU. You should set the source interface MTU to be the network MTU + 306 bytes. So for the standard 1500 MTU network path, the source interface MTU should be 1806, and the cluster control link MTU should be +154, 1960.

For Azure with GWLB, the data interface uses VXLAN encapsulation. In this case, the entire Ethernet datagram is being encapsulated, so the new packet is larger and requires a larger MTU. You should set the cluster control link MTU to be the source interface MTU + 80 bytes.

The following table shows the default values for the cluster control link MTU and the data interface MTU.

Note We do not recommend setting the cluster control link MTU between 2561 and 8362; due to block pool handling, this MTU size is not optimal for system operation.

Table 37: Default MTU

Public Cloud	Cluster Control Link MTU	Data Interface MTU
AWS with GWLB	1980	1826
AWS	1654	1500
Azure with GWLB	1454	1374
Azure	1454	1300
GCP	1554	1400

Guidelines for Threat Defense Virtual Clustering

High Availability

High Availability is not supported with clustering.

IPv6

The cluster control link is only supported using IPv4.

Additional Guidelines

- When significant topology changes occur (such as adding or removing an EtherChannel interface, enabling or disabling an interface on the Firewall Threat Defense or the switch, adding an additional switch to form a VSS or vPC or VNet) you should disable the health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all units, you can re-enable the interface health check feature.
- When adding a node to an existing cluster, or when reloading a node, there will be a temporary, limited packet/connection drop; this is expected behavior. In some cases, the dropped packets can hang your connection; for example, dropping a FIN/ACK packet for an FTP connection will make the FTP client hang. In this case, you need to reestablish the FTP connection.
- Do not power off a node without first disabling clustering on the node.
- For decrypted TLS/SSL connections, the decryption states are not synchronized, and if the connection owner fails, then decrypted connections will be reset. New connections will need to be established to a new node. Connections that are not decrypted (they match a do-not-decrypt rule) are not affected and are replicated correctly.
- Dynamic scaling is not supported.
- If you are using Secure Firewall versions 7.2 or 7.3, Stateful Target Failover is not supported when you deploy the cluster on AWS.
- Perform a global deployment after the completion of each maintenance window.
- Ensure that you do not remove more than one device at a time from the auto scale group (AWS) / instance group (GCP) / scale set (Azure). We also recommend that you run the **cluster disable** command on the device before removing the device from the auto scale group (AWS) / instance group (GCP) / scale set (Azure).
- If you want to disable data nodes and the control node in a cluster, we recommend that you disable the data nodes before disabling the control node. If a control node is disabled while there are other data nodes in the cluster, one of the data nodes has to be promoted to be the control node. Note that the role change could disturb the cluster.
- In the customized day 0 configuration scripts given in this guide, you can change the IP addresses as per your requirement, provide custom interface names, and change the sequence of the CCL-Link interface.
- If you experience CCL instability issues, such as intermittent ping failures, after deploying a Threat Defense Virtual cluster on a cloud platform, we recommend that you address the reasons that are causing CCL instability. Also, you can increase the hold time as a temporary workaround to mitigate CCL instability issues to a certain extent. For more information on how to change the hold time, see Edit Cluster Health Monitor Settings.
- When you are configuring your security firewall rule or security group for the Management Center virtual, you must include both Private and Public IP addresses of the Firewall Threat Defense Virtual in the Source IP address range. Also, ensure to specify the Private and Public IP addresses of the Firewall Management Center Virtual in the security firewall rule or security group of the Firewall Threat Defense Virtual. This is important to ensure proper registration of nodes during clustering deployment.

Defaults for Clustering

- The cLACP system ID is auto-generated, and the system priority is 1 by default.

- The cluster health check feature is enabled by default with the holdtime of 3 seconds. Interface health monitoring is enabled on all interfaces by default.
- The cluster auto-rejoin feature for a failed cluster control link is unlimited attempts every 5 minutes.
- The cluster auto-rejoin feature for a failed data interface is 3 attempts every 5 minutes, with the increasing interval set to 2.
- Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

Deploy the Cluster in AWS

To deploy a cluster in AWS, you can either manually deploy or use CloudFormation templates to deploy a stack. You can use the cluster with AWS Gateway Load Balancer, or with a non-native load-balancer such as the Cisco Cloud Services Router.

AWS Gateway Load Balancer and Geneve Single-Arm Proxy

Note This use case is the only currently supported use case for Geneve interfaces.

The AWS Gateway Load Balancer combines a transparent network gateway and a load balancer that distributes traffic and scales virtual appliances on demand. The Threat Defense Virtual supports the Gateway Load Balancer centralized control plane with a distributed data plane (Gateway Load Balancer endpoint). The following figure shows traffic forwarded to the Gateway Load Balancer from the Gateway Load Balancer endpoint. The Gateway Load Balancer balances traffic among multiple Threat Defense Virtuals, which inspect the traffic before either dropping it or sending it back to the Gateway Load Balancer (U-turn traffic). The Gateway Load Balancer then sends the traffic back to the Gateway Load Balancer endpoint and to the destination.

Note Transport Layer Security (TLS) Server Identity Discovery is not supported with Geneve single-arm setup on AWS.

Figure 267: Geneve Single-Arm Proxy

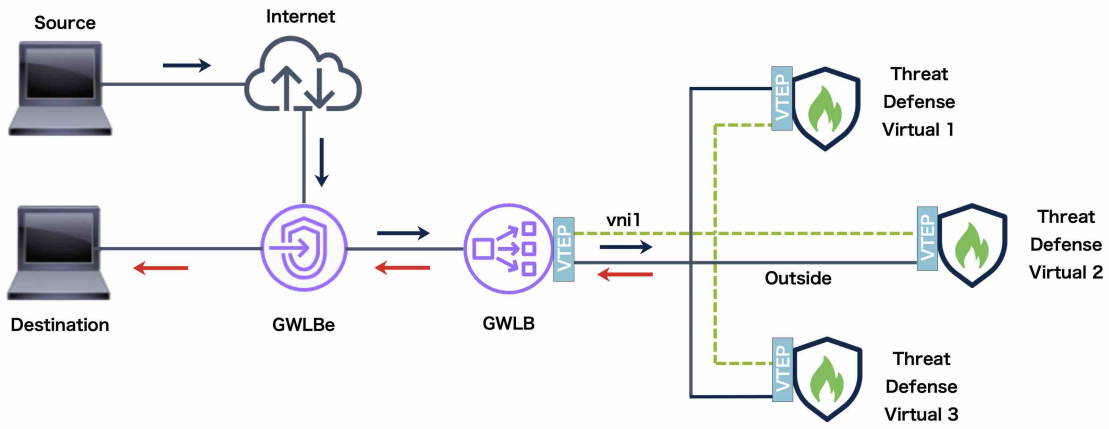

Sample Topologies

Firewall Threat Defense Virtual Clustering with Autoscale in Single and Multiple Availability Zones of an AWS Region

An availability zone is a standalone data center or a set of independent data centers within an AWS region that operate independently. Each zone has its own networking infrastructure, connectivity, and power source ensuring a failure in one zone does not affect others. To improve redundancy and reliability, companies use multiple availability zones in their disaster recovery plans.

Deploying Firewall Threat Defense Virtual across multiple availability zones and configuring clustering with dynamic scaling can significantly enhance the availability and scalability of your infrastructure. In addition, utilizing multiple availability zones in the same region can offer extra redundancy and guarantee high availability in the event of a failure.

You can modify the IP allocation mechanism of Cluster Control Link (CCL) to support both single and multiple availability zone deployments of Firewall Threat Defense Virtual clusters on AWS. The topologies given below depict both inbound and outbound traffic flow in a single and multiple availability zones with autoscaling ability.

Firewall Threat Defense Virtual Clustering with Autoscale in Single Availability Zone

There are two Firewall Threat Defense Virtual instances in the cluster that are connected to a GWLB.

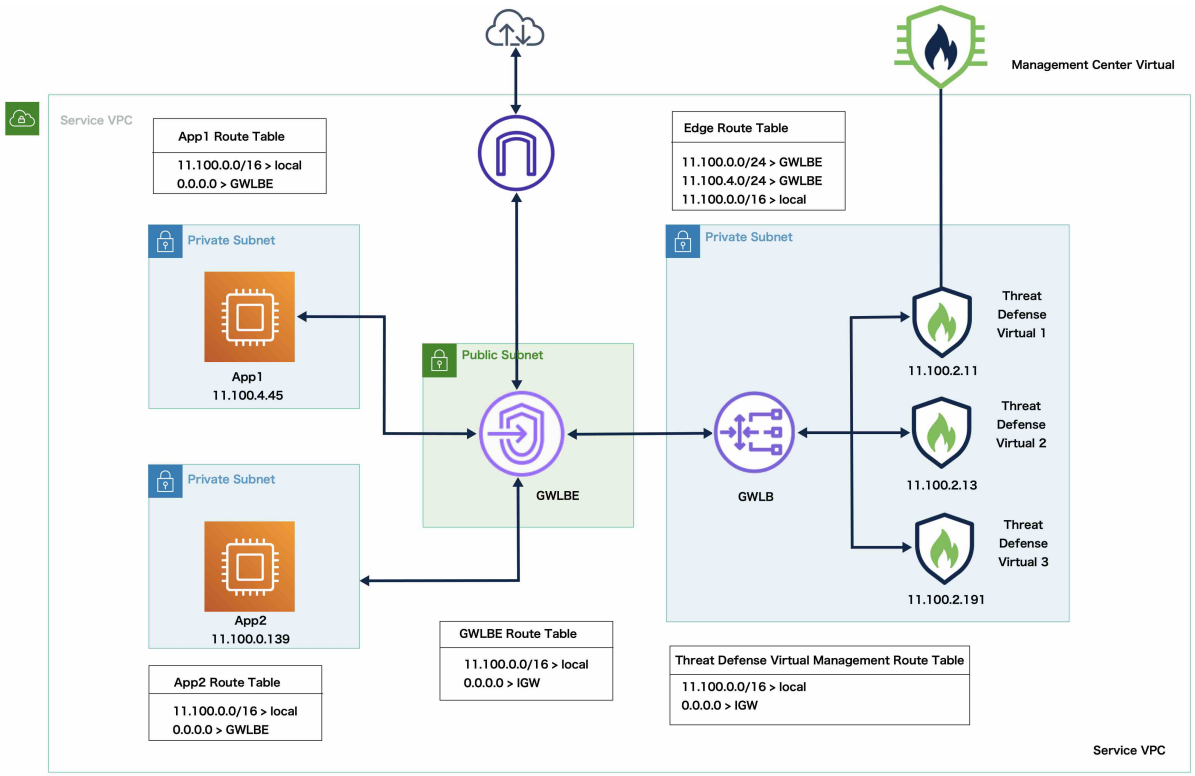

Inbound traffic from the internet goes to the GWLB endpoint, which is then transmits the traffic to the GWLB. Traffic is then forwarded to the Firewall Threat Defense Virtual cluster. After the traffic is inspected by an Firewall Threat Defense Virtual instance in the cluster, it is forwarded to the application VM, App1.

Outbound traffic from App1 is transmitted to the GWLB endpoint > GWLB > TDv > GWLB > GWLB Endpoint, which then sends it out to the internet.

Firewall Threat Defense Virtual Clustering with Autoscale in Multiple Availability Zones

There are two Firewall Threat Defense Virtual instances in the cluster in different availability zones that are connected to a GWLB.

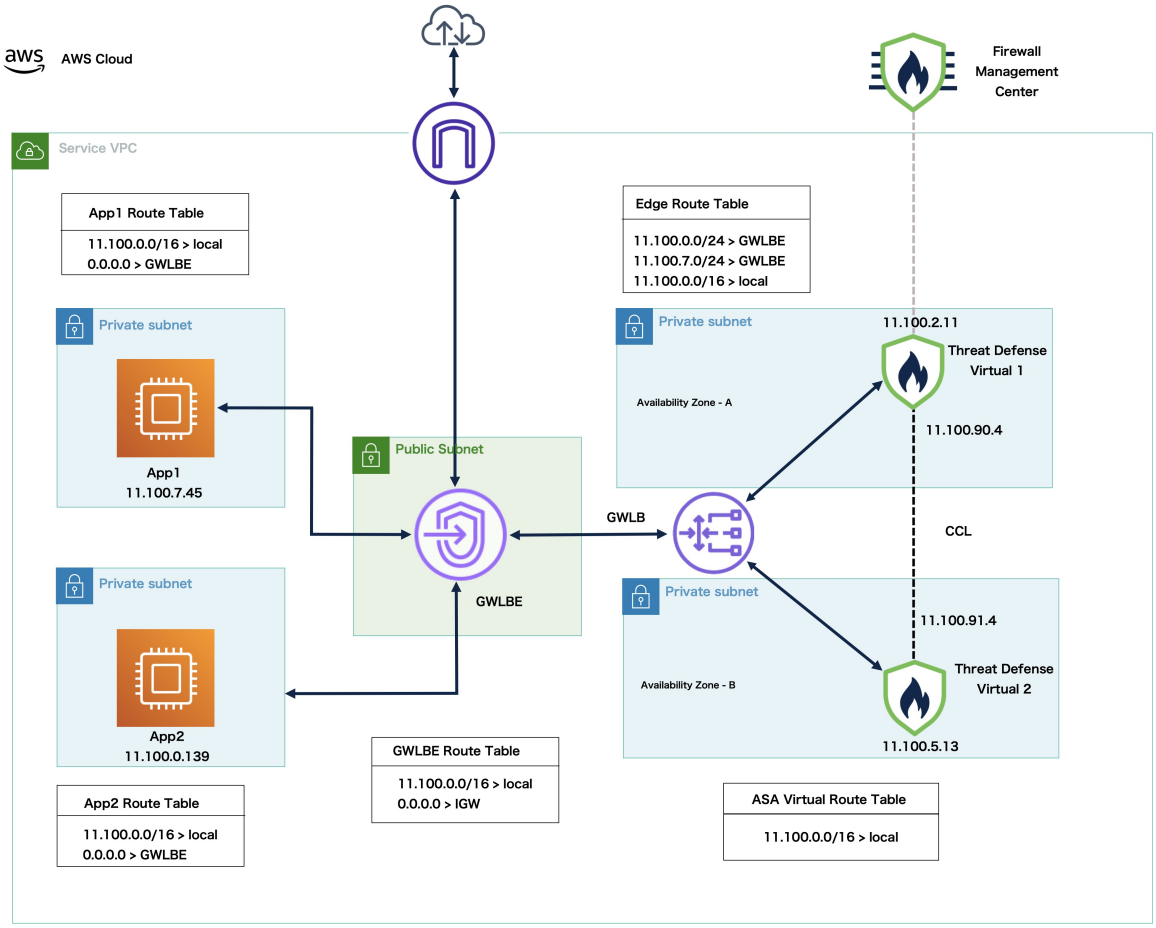

Note Multiple Availability Zone deployment is supported from Firewall Threat Defense Virtual Version 7.6.0 and later.

Inbound traffic from the internet goes to the GWLB endpoint, which then transmits the traffic to the GWLB. Based on the availability zone, the traffic is then routed to the Firewall Threat Defense Virtual cluster. After the traffic is inspected by an Firewall Threat Defense Virtual instance in the cluster, it is forwarded to the application VM, App1.

End-to-End Process for Deploying Threat Defense Virtual Cluster on AWS

Template-based Deployment

The following flowchart illustrates the workflow for template-based deployment of the Threat Defense Virtual cluster on AWS.

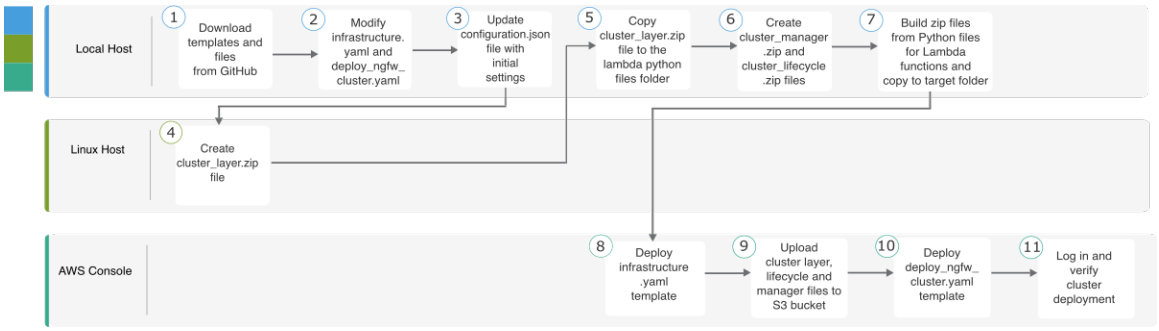

	Workspace	Steps
1	Local Host	Clone the repository from GitHub
2	Local Host	Modify *infrastructure.yaml* and *deploy_ngfw_cluster.yaml* templates.
3	Local Host	Update the *Configuration.json* file with FMC object names.
4	Linux Host	Create *cluster_layer.zip* file.
5	Local Host	Copy *cluster_layer.zip* file to the Lambda python files folder.
6	Local Host	Create *cluster_manager.zip*, *custom_metrics_publisher.zip*, and *cluster_lifecycle.zip* files.
7	Local Host	Build zip files from Python files for Lambda functions and copy to target folder.
8	AWS Console	Deploy *infrastructure.yaml* template.
9	AWS Console	Upload *cluster_layer.zip*, *cluster_lifecycle.zip*, *custom_metrics_publisher.zip*, and *cluster_manager.zip* to the S3 bucket.
10	AWS Console	Deploy *deploy_ngfw_cluster.yaml* template.
11	AWS Console	Log in and verify cluster deployment.

Manual Deployment

The following flowchart illustrates the workflow for manual deployment of the Threat Defense Virtual cluster on AWS.

	Workspace	Steps
1	Local Host	Create day 0 configuration script.
2	AWS Console	Deploy Threat Defense Virtual instance.
3	AWS Console	Attach interfaces to instance.
4	AWS Console	Verify if nodes have joined cluster.
5	AWS Console	Create target group and GWLB; attach target group to the GWLB.
6	AWS Console	Register instances with the target group using data interface IP.
7	Management Center	Register control node.

Templates

The templates given below are available in GitHub. The parameter values are self-explanatory with the parameter names, default values, allowed values, and description, given in the template.

- infrastructure.yaml – Template for infrastructure deployment.
- deploy_ngfw_cluster.yaml – Template for cluster deployment.

Note Ensure that you check the list of supported AWS instance types before deploying cluster nodes. This list is found in the *deploy_ngfw_cluster.yaml* template, under allowed values for the parameter InstanceType.

Deploy the Stack in AWS Using a CloudFormation Template

Deploy the stack in AWS using the customized CloudFormation template.

Before you begin

- You need a Amazon Linux virtual machine with Python 3.
- To allow the cluster to auto-register with the Firewall Management Center, you need to create *two* users with administrative privileges on the Firewall Management Center that can use the REST API. See the Cisco Secure Firewall Management Center Administration Guide.
- Add an access policy in the Firewall Management Center that matches the name of the policy that you specified in `Configuration.json`.

Procedure

Step 1 Prepare the template.

a) Clone the GitHub repository to your local folder. See https://github.com/CiscoDevNet/cisco-ftdv/tree/master/cluster/aws.

b) Modify **infrastructure.yaml** and **deploy_ngfw_cluster.yaml** with the required parameters.

c) Modify **cluster/aws/lambda-python-files/Configuration.json** with initial settings.

For example:

```
{
  "licenseCaps": ["BASE", "MALWARE", "THREAT"],
  "performanceTier": "FTDv50",
  "fmcIpforDeviceReg": "DONTRESOLVE",
  "RegistrationId": "cisco",
  "NatId": "cisco",
  "fmcAccessPolicyName": "AWS-ACL"
}
```

- Keep the fmcIpforDeviceReg setting as DONTRESOLVE.
- The fmcAccessPolicyName needs to match an access policy on the Firewall Management Center.

Note
FTDv5 and FTDv10 tiers are not supported.

d) Create a file named **cluster_layer.zip** to provide essential Python libraries to Lambda functions.

We recommend to use the Amazon Linux with Python 3.9 installed to create the **cluster_layer.zip** file.

Note
If you need an Amazon Linux environment, you can create an EC2 instance using Amazon Linux 2023 AMI or use AWS Cloudshell, which runs the latest version of Amazon Linux.

For creating the cluster-layer.zip file, you need to first create **requirements.txt** file that consists of the python library package details and then run the shell script.

1. Create the **requirements.txt** file by specifying the python package details.

The following is the sample package details that you provide in the **requirements.txt** file:

```
$ cat requirements.txt
pycryptodome
paramiko
requests
scp
jsonschema
cffi
zipp
importlib-metadata
```

2. Run the following shell script to create **cluster_layer.zip** file.

```
$ pip3 install --platform manylinux2014_x86_64
--target=./python/lib/python3.9/site-packages
--implementation cp --python-version 3.9 --only-binary=:all:
--upgrade -r requirements.txt
$ zip -r cluster_layer.zip ./python
```

Note
If you encounter a dependency conflict error during installation, such as urllib3 or cryptography, it is recommended that you include the conflicting packages along with their recommended versions in the **requirements.txt** file. After that, you can run the installation again to resolve the conflict.

e) Copy the resulting **cluster_layer.zip** file to the lambda python files folder - `cluster/aws/lambda-python-files`.

f) Create the **cluster_layer.zip**, **custom_metrics_publisher.zip**, **cluster_manger.zip** and **lifecycle_ftdv.zip** files.

A **make.py** file can be found in the cloned repository (`cluster/aws/make.py`). This will zip the python files into a Zip file and copy to a target folder.

python3 make.py build

Note
If you are using a private IP address for the Management Center Virtual registration, then make sure that you set USE_PUBLIC_IP_FOR_FMC_CONN to False in the `cisco-ftdv/cluster/aws/lambda-python-files/constant.py` file.

Step 2 Deploy **infrastructure.yaml** and note the output values for cluster deployment. Before deploying the infrastructure stack, it is important to identify the AWS region and the availability zones that will be used. Each AWS region has a different set of availability zones and VPC infrastructure, therefore it is essential to select the correct region and availability zones for your deployment.

a) On the AWS Console, go to **CloudFormation** and click **Create stack**; select **With new resources(standard)**.

b) Select **Upload a template file**, click **Choose file**, and select **infrastructure.yaml** from the target folder.

c) Click **Next** and provide the required information.

Parameter	Allowed Values/Type	Description
ClusterName	String	Enter unique Cluster name.
ClusterNumber	Number	Enter unique Cluster number.
VpcCidr	String	Enter the CIDR block for a new VPC

Parameter	Allowed Values/Type	Description
NoOfAZs	Number	Select 2 or 3 Availability Zones (AZs) for releases 7.6.0 and above; for lower releases, Select 1AZ. Management, Inside, Outside and CCL subnets will be distributed across the chosen AZs accordingly.
ListOfAZs	List	Select Availability Zones (Count should match with Number of Availability Zones)
MgmtSubnetNames	CommaDelimitedList	Management subnets name (With Internet GW as Route)
MgmtSubnetCidrs	CommaDelimitedList	Management subnets Cidr list
InsideSubnetNames	CommaDelimitedList	Inside subnets name (With Private Route)
InsideSubnetCidrs	CommaDelimitedList	Inside subnets Cidr list
CCLSubnetNames	CommaDelimitedList	Enter CCL subnet name
CCLSubnetCidrs	CommaDelimitedList	Enter CCL subnet CIDR
LambdaAZs	List	Select 2 Availability Zones for Lambda
LambdaSubnetNames	CommaDelimitedList	Enter Lambda Subnets name (With NAT GW as Route), for Lambda Functions
LambdaSubnetCidrs	CommaDelimitedList	Enter Lambda Subnet CIDRs

d) Enter a unique **Cluster Name** and **Cluster Number** for the cluster.
e) Select the availability zone from the **Availability Zone** list. This field lists only availability zones based on the AWS region that you select for deploying the infrastructure stack using the ClusterFormation template.
f) Click **Next**, then **Create stack**.
g) After the deployment is complete, go to **Outputs** and note the S3 **BucketName**.

Figure 268: Output of infrastructure.yaml

Outputs (13)

Search outputs

Key	Value	Description	Export name
BucketName	maz-cls-infra-s3bucketcluster-whpzbytieb2h	Name of the Amazon S3 bucket	-
BucketUrl	http://maz-cls-infra-s3bucketcluster-whpzbytieb2h.s3-website-us-east-1.amazonaws.com	URL of S3 Bucket Static Website	-
CCLSubnetIds	subnet-0bc04e2cc9e53e5c0,subnet-0d7d046a0fca25615,subnet-03ef42bf527551569	List of CCL subnet IDs (comma seperated)	-
EIPforNATgw	3.218.44.132	EIP reserved for NAT GW	-
FmcInstanceSGID	sg-076880aa64df2db5c	Security Group ID for FMC if user would like to launch in this VPC itself	-
InInterfaceSGId	sg-06ed933d6624fe51b	Security Group ID for Inside Interfaces	-
InsideSubnetIds	subnet-03d12cab8ee0eafff,subnet-0be9158b0970aebab,subnet-0b53c96fceb7c1f4d	List of Inside subnet IDs (comma seperated)	-
InstanceSGId	sg-0680b74be473186aa	Security Group ID for Instances Management Interface	-
LambdaSecurityGroupId	sg-057da2a9954e0d204	Security Group ID for Lambda Functions	-
LambdaSubnetIds	subnet-03439803d989e6bdf,subnet-087488a9d6ffc95cd	List of lambda subnet IDs (comma seperated)	-
ListOfAZs	us-east-1a,us-east-1b,us-east-1c	Availability zones for NGFWv instances	-
MgmtSubnetIds	subnet-06f0bbbd3f207a504,subnet-0c339dc43688cddc9,subnet-0a67629632a655de7	List of Mangement subnet IDs (comma seperated)	-
VpcName	vpc-09c2b0ad995e2fb24	Name of the VPC created	-

Step 3 Upload **cluster_layer.zip**, **cluster_manager.zip**, **custom_metrics_publisher.zip**, and **cluster_lifecycle.zip** to the S3 bucket created by **infrastructure.yaml**.

Figure 269: S3 Bucket

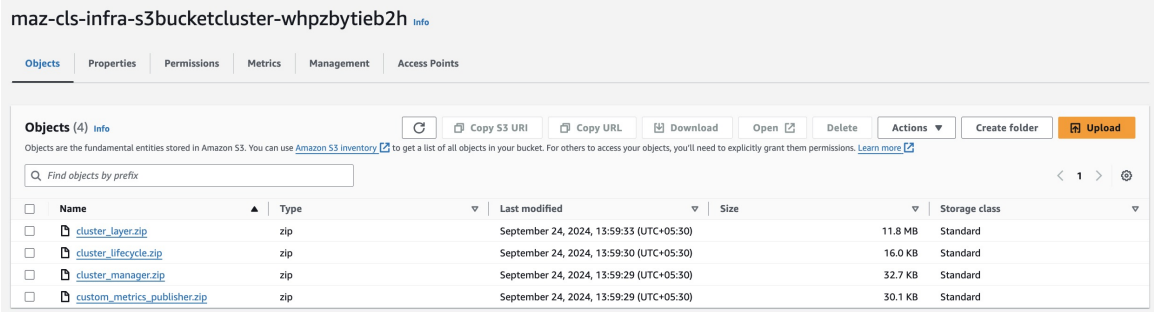

Note
Make sure that the Elastic IP address of the Lambda NAT Gateway is added to the security group associated with the Management Center Virtual.

Step 4 Deploy **deploy_ngfw_cluster.yaml**.

a) Go to **CloudFormation** and click on **Create stack**; select **With new resources(standard)**.
b) Select **Upload a template file**, click **Choose file**, and select **deploy_ngfw_cluster.yaml** from the target folder.
c) Click **Next** and provide the required information.
d) Provide the following cluster and infrastructure configuration information.

Parameter	Allowed Values/Type	Description
Cluster Configuration		
ClusterGrpNamePrefix	String	This is the cluster name Prefix. The cluster number will be added as a suffix.
ClusterNumber	String	This is the cluster number. This will be suffixed to the cluster name (msa-ftdv-infra). For example, if this value is **1**, the group name will be *msa-ftdv-infra-1*. It should be at least 1 digit, but not more than 3 digits. Default: 1.
ClusterSize	Numbers	This is the total number of Firewall Threat Defense Virtual nodes in a cluster. Minimum: 1 Maximum:16
Infrastructure Details		
NoOfAZs	String	This is the total number of availability zones into which Firewall Threat Defense Virtual is deployed. (The number of availability zones varies from a Minimum 1 to Maximum 3 depending on a region). The subnet will be created in these availability zones.

Parameter	Allowed Values/Type	Description
		The availability zones available in this list is based on the region selected for deploying the cluster. **Note** Management, Inside, and Cluster Control Link (CCL) subnets are created across three availability zones based on this parameters.
AZ	String	The availability zone list is based on the region you plan to deploy. In Availability Zone list, select the valid availability zone (1 availability zone or 2 availability zones or 3 availability zones). Count should match with the value of Number of Availability Zones parameter.
NotifyEmailID	String	Email address to which cluster events email will be sent. You must accept a subscription email request to receive this email notification. Example: admin@company.com
VpcId	String	The VPC ID for the cluster group. Type: AWS::EC2::VPC::Id
S3BktName	String	The S3 Bucket that contains the uploaded Lambda zip files. You must specify correct bucket name.
MgmtSubnetIds	List	Enter only *one* subnet per availability zone. If you select multiple subnets from a same availability zone, then selecting an incorrect subnet may cause issues while deploying the Firewall Threat Defense Virtual instances. Type: List<AWS::EC2::Subnet::Id>
InsideSubnetIds	List	Enter at least *one* subnet per availability zone. If multiple subnets from the same Availability Zone are selected, then selecting an incorrect subnet may cause issues while deploying the Firewall Threat Defense Virtual instances. Type: List<AWS::EC2::Subnet::Id>
LambdaSubnets	List	Enter at least *two* subnet for the Lambda functions. The *two* subnets you enter must have a NAT gateway to enable the Lambda functions to communicate with AWS services, which are public DNS. Type: List<AWS::EC2::Subnet::Id>

Parameter	Allowed Values/Type	Description
CCLSubnetIds	String	Enter at least *one* subnet per availability zone. If multiple subnets from the same Availability Zone are selected, then selecting an incorrect subnet may cause issues while deploying the Firewall Threat Defense Virtual instances. Type: List<AWS::EC2::Subnet::Id>
CCLSubnetRanges	String	Enter IP addresses range of CCL subnets for different availability zones. Exclude first 4 reserved IP addresses. IP address pool for Cluster Control Link (CCL). IP address is allocated to the CCL interfaces of the Firewall Threat Defense Virtual instance from CCL IP address pool.
MgmtInterfaceSG	List	Select security group ID for the Firewall Threat Defense Virtual instances. Type: List<AWS::EC2::SecurityGroup::Id>
InsideInterfaceSG	List	Select security group ID for the inside interface of Firewall Threat Defense Virtual instances. Type: List<AWS::EC2::SecurityGroup::Id>
LambdaSG	List	Select a security group for the Lambda functions. Ensure outbound connections is set to **ANYWHERE**. Type: List<AWS::EC2::SecurityGroup::Id>
CCLInterfaceSG	List	Select a security group ID for CCL interface of the Firewall Threat Defense Virtual instances.
GWLB Configuration		
DeployGWLBE	String	Click **Yes** to deploy the GWLB endpoint. By default, the value is set to **No**.
VpcIdLBE	String	Enter VPC to deploy Gateway Load Balancer Endpoint. **Note** Do not enter any value in this field if you are not deploying the GWLB endpoint.
GWLBESubnetId	String	Enter only one subnet ID. **Note** Do not enter any value in this field if you are not deploying the GWLB endpoint.

Parameter	Allowed Values/Type	Description
		Ensure that the subnet belongs to the correct VPC, and the availability zones that you have specified.
TargetFailover	String	Enable Target Failover support when a target fails or deregisters. (By default, the value of this parameter is set to **rebalance**). • **no_rebalance**: Directs existing flows to failed targets and new flows to healthy targets, ensuring backward compatibility. • **rebalance**: Redistributes existing flows while ensuring that new flows go to healthy targets. *rebalance* is supported from Firewall Threat Defense Virtual Version 7.4.1 and later.
TgHealthPort	String	Enter Health Check Port for GWLB. **Note** By default, this port must not be used for traffic. Ensure the value you provide is a valid TCP port. Default: 8080
Cisco NGFWv Instance Configuration		
InstanceType	String	Cisco Firewall Threat Defense Virtual EC2 instance type. Ensure that the AWS Region supports Instance Type you select. By default, **c5.xlarge** is selected.
LicenseType	String	Choose Cisco Firewall Threat Defense Virtual EC2 instance license type. Ensure that the AMI ID that you enter in **AMI-ID** parameter is of the same licensing type. By default, **BYOL** is selected.
AssignPublicIP	String	Set the value as **true** to assign a public IP address for Firewall Threat Defense Virtual from the AWS IP address pool.
AmiID	String	Choose the correct AMI ID as per the region, version, and license type (BYOL or PAYG). Firewall Threat Defense Virtual 7.2 and later support clustering, and Firewall Threat Defense Virtual Version 7.6 and later support the autoscaling and multiple availability zone enhancements.

Parameter	Allowed Values/Type	Description
		Type: AWS::EC2::Image::Id
ngfwPassword	String	Firewall Threat Defense Virtual instance password. All Firewall Threat Defense Virtual instances come up with a default password, which is in the Userdata field of the Launch Template (Cluster Group). The password is activated after Firewall Threat Defense Virtual is accessible. Minimum length must be 8 characters. The password can either be a plain text password or a KMS encrypted password.
KmsArn	String	Enter ARN of an existing KMS (AWS KMS key to encrypt at rest). If you specify a value in this field, then the Firewall Threat Defense Virtual instance's ***admin*** password must be an encrypted password. Example of generating an encrypted password: "*aws kms encrypt --key-id <KMS ARN> --plaintext <password>* " The password encryption must be done using only the specified ARN.
FMC Automation Configuration		
fmcDeviceGrpName	String	Enter a unique name for the cluster group in management center.
fmcPublishMetrics	String	Select **`true`** to create a Lambda Function to poll management center and publish specific device group metrics to AWS CloudWatch. Allowed values: • true • false By default, the value is set to **true**.
fmcMetricsUsername	String	Enter a unique internal user name for polling memory metrics from management center. The user must have privileges of **Network Admin** and **Maintenance User** or more .
fmcMetricsPassword	String	Enter the password.

Parameter	Allowed Values/Type	Description
		If you have mentioned KMS Master Key ARN parameter, ensure to provide an encrypted password. Ensure to enter the correct password because entering incorrect password may result in failure of metrics collection.
fmcServer	String	The IP address can be an external IP address or the IP address reachable in Firewall Threat Defense Virtual management subnet in the VPC. Minimum length: 7 Maximum length:15
fmcOperationsUsername	String	Provide a unique internal user name for Firewall Management Center Virtual for CloudWatch. The user must have **Administrator** privileges.
fmcOperationsPassword	String	Enter the password. If you have mentioned KMS Master Key ARN parameter, ensure to provide an encrypted password.
Scaling Configuration		
CpuThresholds	CommaDelimitedList	(Optional) Specifying non-zero lower and upper thresholds will create scale policies. If (0,0) is selected, no CPU scaling alarm or policies will be created. Evaluation points and data points are at default or recommended values. By default, **Autoscale** is enabled in this template. Autoscale can be disabled after deployment.
MemoryThresholds	CommaDelimitedList	Specifying non-zero lower and upper threshold will create scale policies. If (0,0) is selected, no memory scaling alarm or policies will be created. Evaluation points and data points are at default or recommended values.

e) Click **Next**.
f) Click to acknowledge all the AWS CloudFormation options.
g) Click **Submit** to deploy the cluster.
h) Click **Next**, then **Create stack**.

The Lambda functions manage the rest of the process, and the Firewall Threat Defense Virtuals will automatically register with the Firewall Management Center.

Figure 270: Deployed Resources

The status changes from **CREATE_IN_PROGRESS** to **CREATE COMPLETE** indicating successful deployment.

Step 5 Verify the cluster deployment by logging into any one of the nodes and using the **show cluster info** command.

Figure 271: Cluster Nodes

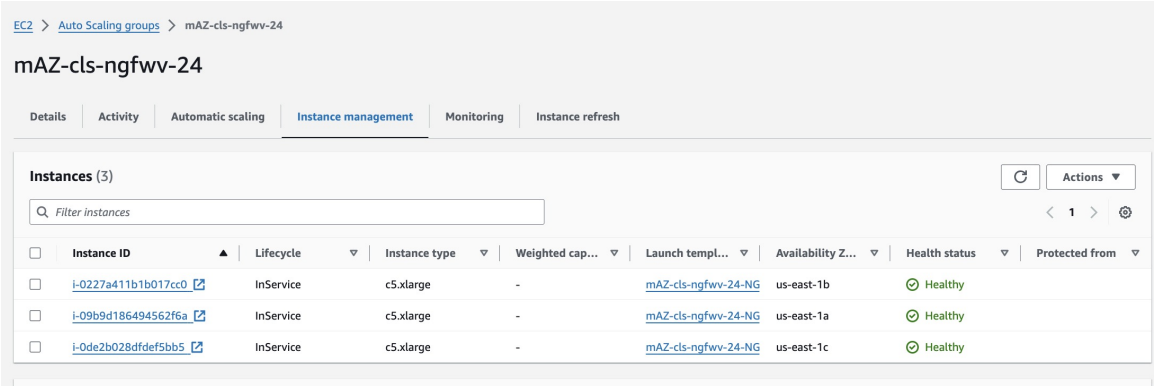

Figure 272: show cluster info

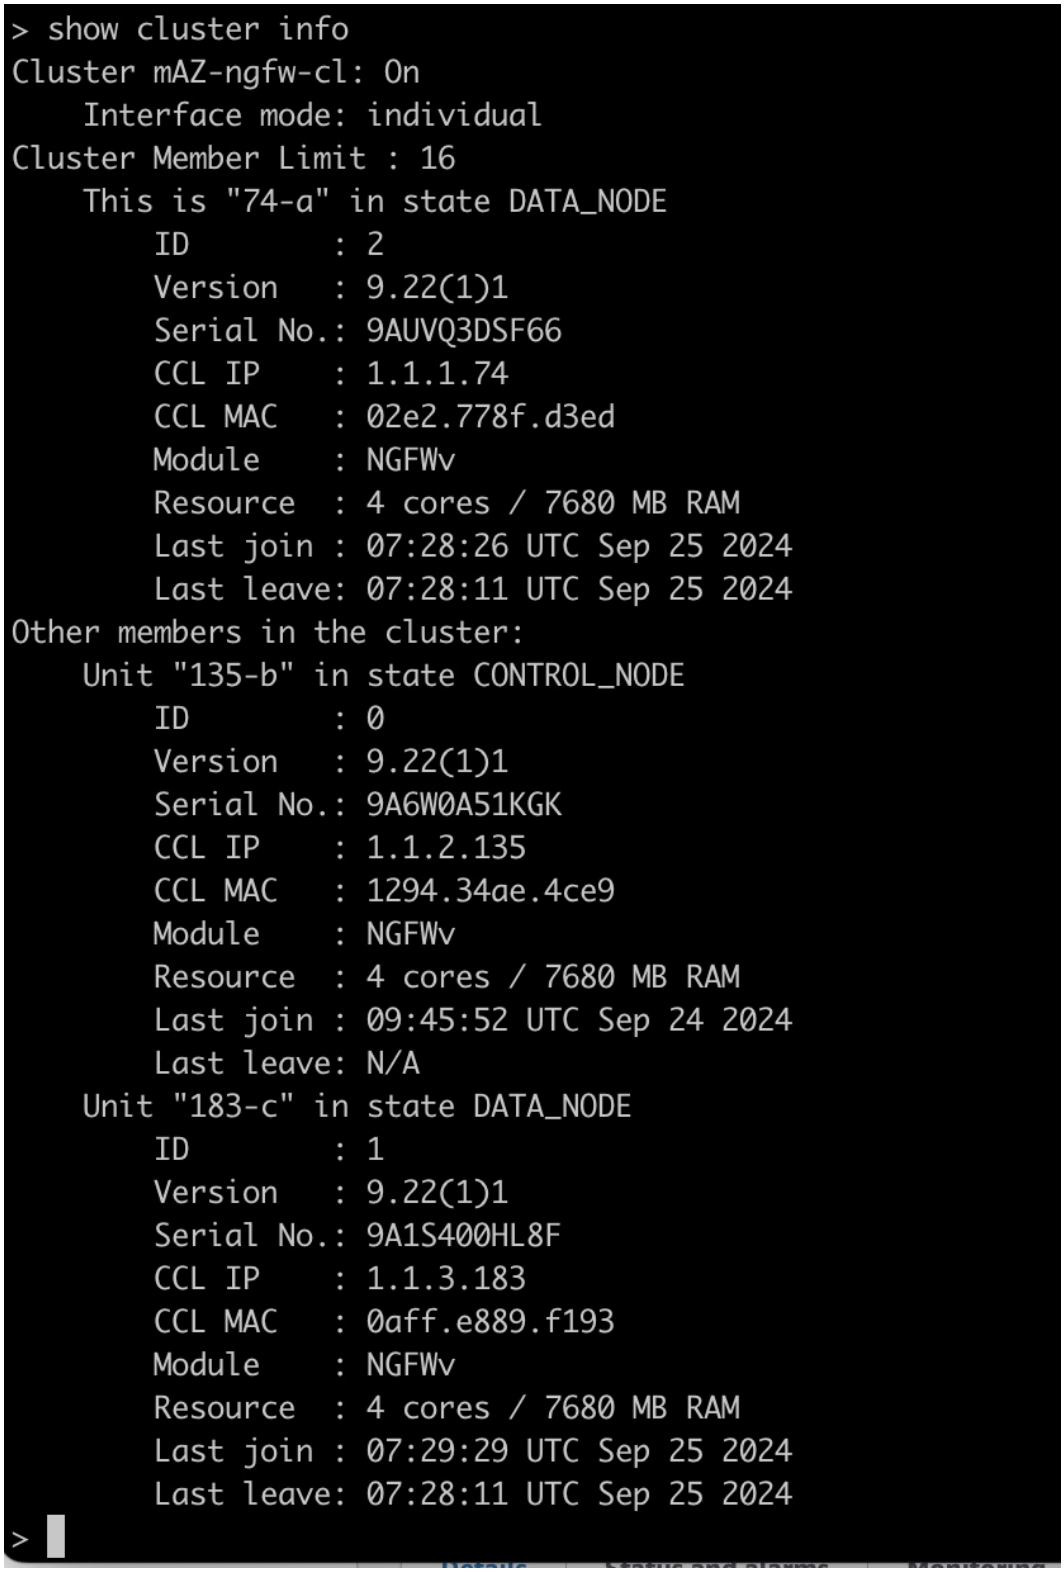

Autoscale Parameter Configuration

After the deployment is completed, you must specify **Minimum**, **Maximum**, and **Desired** capacity of the Firewall Threat Defense Virtual Autoscale group. You must verify the Autoscale functionality.

Procedure

Step 1 From the AWS console, choose **Services** > **EC2** > **Auto Scaling groups** > **Created ClusterAutoscale group**.

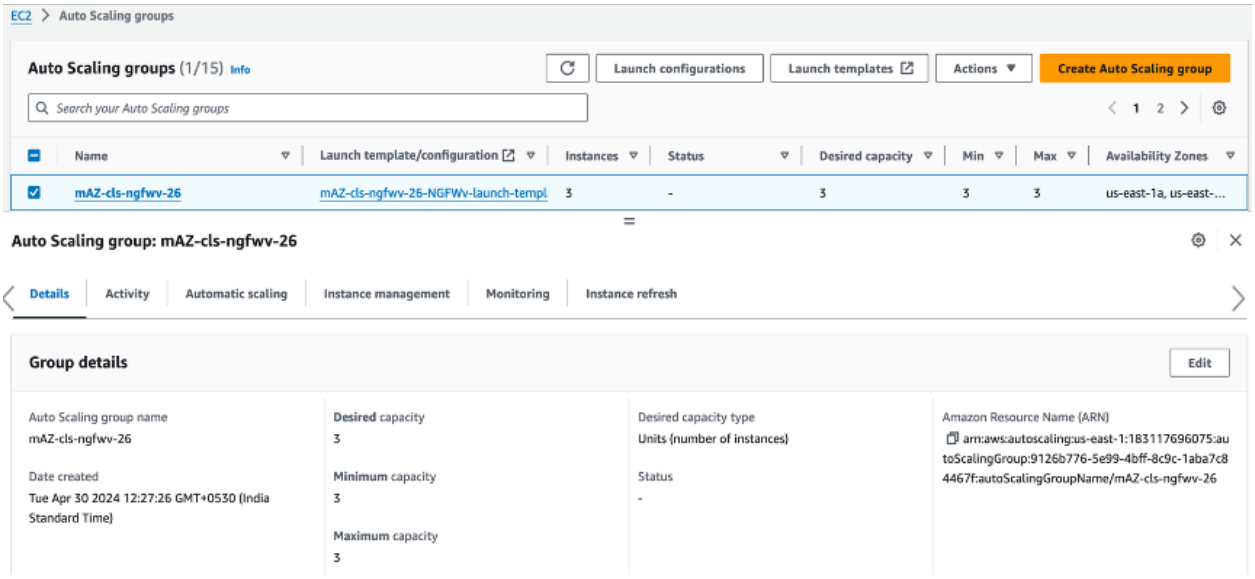

Step 2 Select the autoscale group check box.

Step 3 Click **Actions** to edit the autoscaling group capacity.

Step 4 Configure **Desired capacity**, and then set the **Scaling limits** capacity.

Step 5 Check if the CPU and Memory metric data is available and whether scaling is occurring as expected in AWS Cloudwatch alarms.

Configure IMDSv2 Required Mode in Firewall Threat Defense Virtual Clustering by Updating Stack

You can configure the IMDSv2 Required mode for the Firewall Threat Defense Virtual autoscale group instances that are already deployed on the AWS.

Before you begin

IMDSv2 Required mode is only supported by Firewall Threat Defense Virtual version 7.6 and later. You must ensure that your existing instances version is compatible (upgraded to version 7.6) with IMDSv2 mode before configuring the IMDSv2 mode for your deployment.

Procedure

Step 1 On the AWS Console, go to **CloudFormation** and click **Stacks**.
Step 2 Select the stack of the intially deployed clustering instances.
Step 3 Click **Update**.
Step 4 On the **Update stack** page, click **Replace existing template**.
Step 5 Under **Specify template** section, click **Upload a template file**.
Step 6 Choose and upload the template which support IMDSv2.
Step 7 Provide values for the input parameters in the template.
Step 8 Update the stack.

Deploy the Cluster in AWS Manually

To deploy the cluster manually, prepare the day 0 configuration, deploy each node, and then add the control node to the Firewall Management Center.

Create the Day0 Configuration for AWS

You can use either a fixed configuration or a customized configuration. We recommend using the fixed configuration.

Create the Day0 Configuration With a Fixed Configuration for AWS

The fixed configuration will auto-generate the cluster bootstrap configuration.

Single Availability Zone - Day0 Configuration with a fixed configuration for AWS

```
{
    "AdminPassword": "password",
    "Hostname": "hostname",
    "FirewallMode": "Routed",
    "ManageLocally": "No",
    "Cluster": {
        "CclSubnetRange": "ip_address_start ip_address_end",
        "ClusterGroupName": "cluster_name",
        [For Gateway Load Balancer] "Geneve": "{Yes | No}",
        [For Gateway Load Balancer] "HealthProbePort": "port"
    }
}
```

For example:

```
{
 "AdminPassword": "Sup3rnatural",
 "Hostname": "ciscoftdv",
 "FirewallMode": "Routed",
 "ManageLocally": "No",
 "Cluster": {
             "CclSubnetRange": "10.5.90.4 10.5.90.30",
```

```
  "ClusterGroupName": "ftdv-cluster",
  "Geneve": "Yes",
  "HealthProbePort": "7777"
 }
}
```

Multiple Availability Zone - Day0 Configuration with a fixed configuration for AWS

```
{
    "AdminPassword": "password",
    "Hostname": "hostname",
    "FirewallMode": "Routed",
    "ManageLocally": "No",
    "Cluster": {
    "CclSubnetRange":[
    "ip_address_start_AZ1 ip_address_end_AZ1",
    "ip_address_start_AZ2 ip_address_end_AZ2",
    "ip_address_start_AZ3 ip_address_end_AZ3"
     ],
   "ClusterGroupName": "cluster_name",
        [For Gateway Load Balancer] "Geneve": "{Yes | No}",
        [For Gateway Load Balancer] "HealthProbePort": "port"
    }
}
```

For example: **Two Availability Zone**

```
{
 "AdminPassword": "Sup4rnatural",
 "Hostname": "ftdvcluster",
 "FirewallMode": "Routed",
 "ManageLocally": "No",
 "Cluster": {
        "CclSubnetRange": [
            "10.5.90.4  10.5.90.30",
            "10.5.91.4  10.5.91.30"
        ],
  "ClusterGroupName": "ftdv-cluster",
  "Geneve": "Yes",
  "HealthProbePort": "8080"
 }
}
```

For example: **Three Availability Zone**

```
{
 "AdminPassword": "Sup4rnatural",
 "Hostname": "ftdvcluster",
 "FirewallMode": "Routed",
 "ManageLocally": "No",
 "Cluster": {
        "CclSubnetRange": [
            "10.5.90.4  10.5.90.30",
            "10.5.91.4  10.5.91.30",
            "10.5.92.4  10.5.92.30"
        ],
  "ClusterGroupName": "ftdv-cluster",
  "Geneve": "Yes",
  "HealthProbePort": "8080"
 }
}
```

For the **CclSubnetRange** variable, specify a range of IP addresses starting from x.x.x.4. Ensure that you have at least 16 available IP addresses for clustering. Some examples of start (`ip_address_start`) and end (`ip_address_end`) IP addresses given below.

Table 38: Examples of Start and End IP addresses

CIDR	Start IP Address	End IP Address
10.1.1.0/27	10.1.1.4	10.1.1.30
10.1.1.32/27	10.1.1.36	10.1.1.62
10.1.1.64/27	10.1.1.68	10.1.1.94
10.1.1.96/27	10.1.1.100	10.1.1.126
10.1.1.128/27	10.1.1.132	10.1.1.158
10.1.1.160/27	10.1.1.164	10.1.1.190
10.1.1.192/27	10.1.1.196	10.1.1.222
10.1.1.224/27	10.1.1.228	10.1.1.254
10.1.1.0/24	10.1.1.4	10.1.1.254

Deploy Cluster Nodes

Deploy the cluster nodes so they form a cluster.

Procedure

Step 1 Deploy the Threat Defense Virtual instance by using the cluster day 0 configuration with the required number of interfaces - four interfaces if you are using Gateway Load Balancer (GWLB), or five interfaces if you are using non-native load balancer. To do this, in the **Configure Instance Details** > **Advanced Details** section, paste the cluster day 0 configuration.

Note
Ensure that you attach interfaces to the instances in the order given below.

- AWS Gateway Load Balancer - four interfaces - management, diagnostic, inside, and cluster control link.
- Non-native load balancers - five interfaces - management, diagnostic, inside, outside, and cluster control link.

For more information on deploying Threat Defense Virtual on AWS, see Deploy the Threat Defense Virtual on AWS.

Step 2 Repeat Step 1 to deploy the required number of additional nodes.

Step 3 Use the **show cluster info** command on the Threat Defense Virtual console to verify if all nodes have successfully joined the cluster.

Step 4 Configure the AWS Gateway Load Balancer.

a) Create a target group and GWLB.

b) Attach the target group to the GWLB.

Note
Ensure that you configure the GWLB to use the correct security group, listener configuration, and health check settings.

c) Register the data interface (inside interface) with the Target Group using IP addresses.

For more information, see Create a Gateway Load Balancer.

Step 5 Add the control node to the Management Center. See Add the Cluster to the Management Center (Manual Deployment), on page 662.

Configure Target Failover for Secure Firewall Threat Defense Virtual Clustering with GWLB in AWS

Threat Defense Virtual clustering in AWS utilizes the Gateway Load Balancer (GWLB) to balance and forward network packets for inspection to a designated Threat Defense Virtual node. The GWLB is designed to continue sending network packets to the target node in the event of a failover or deregistration of that node.

The Target Failover feature in AWS enables GWLB to redirect network packets to a healthy target node in the event of node deregistration during planned maintenance or a target node failure. It takes advantage of the cluster's stateful failover.

In AWS, you can configure Target Failover through the AWS Elastic Load Balancing (ELB) API or AWS console.

Note If a target node fails while the GWLB routes traffic using certain protocols such as SSH, SCP, CURL, and so on, then there may be a delay in redirecting traffic to a healthy target. This delay is due to rebalancing and rerouting of traffic flow.

In AWS, you can configure Target Failover through the AWS ELB API or AWS console.

- AWS API - In the AWS ELB API - *modify-target-group-attributes* you can define the flow handling behavior by modifying the following two new parameters.
  - target_failover.on_unhealthy - It defines how the GWLB handles the network flow when the target becomes unhealthy.
  - target_failover.on_deregistration - It defines how the GWLB handles the network flow when the target is deregistered.

  The following command shows the sample API parameter configuration of defining these two parameters.

```
aws elbv2 modify-target-group-attributes \
--target-group-arn arn:aws:elasticloadbalancing:…/my-targets/73e2d6bc24d8a067 \
--attributes \
Key=target_failover.on_unhealthy, Value=rebalance[no_rebalance] \
Key=target_failover.on_deregistration, Value=rebalance[no_rebalance]
```

For more information, refer TargetGroupAttribute in the AWS documentation.

- AWS Console – In the EC2 console, you can enable the Target Failover option on the Target Group page by configuring the following options.
  - Edit Target Groups Attributes
  - Enable Target Failover
  - Verify Rebalance Flows

For more information about how to enable Target Failover, see Enable Target Failover for Secure Firewall Threat Defense Virtual Clustering in AWS, on page 609.

Enable Target Failover for Secure Firewall Threat Defense Virtual Clustering in AWS

The data interface of Firewall Threat Defense Virtual is registered to a target group of GWLB in AWS. In the Firewall Threat Defense Virtual clustering, each instance is associated with a Target Group. The GWLB load balances and sends the traffic to this healthy instance identified or registered as a target node in the target group.

Before you begin

You must have deployed the cluster in AWS either by manual method or using CloudFormation templates.

If you are deploying a cluster using a CloudFormation template, you can also enable the **Target Failover** parameter by assigning the **rebalance** attribute that is available under **GWLB Configuration** section of the cluster deployment file, `deploy_ftdv_clustering.yaml`. In the template, by default, the value is set to **rebalance** for this parameter. However, the default value for this parameter is set to **no_rebalance** on the AWS console.

Where,

- **no_rebalance** - GWLB continues to send the network flow to the failed or deregistered target.
- **rebalance** - GWLB sends the network flow to another healthy target when the existing target is failed or deregistered.

For information on deploying stack in AWS, see:

- Deploy the Cluster in AWS Manually
- Deploy the Stack in AWS Using a CloudFormation Template

Procedure

Step 1 On the AWS Console, go to **Services** > **EC2**

Step 2 Click **Target Groups** to view the target groups page.

Step 3 Select the target group to which the Firewall Threat Defense Virtual data interface IPs are registered. The target group details page is displayed, where you can enable the Target failover attributes.

Step 4 Go to the **Attributes** menu.

Step 5 Click **Edit** to edit the attributes.

Step 6 Toggle the **Rebalance flows** slider button to the right to enable target failover to configure GWLB to rebalance and forward the existing network packets to a healthy target node in the event of target failover or deregistration.

Deploy the Cluster in Azure

You can use the cluster with the Azure Gateway Load Balancer (GWLB), or with a non-native load-balancer. To deploy a cluster in Azure, use Azure Resource Manager (ARM) templates to deploy a Virtual Machine Scale Set.

Sample Topology for GWLB-based Cluster Deployment

Figure 273: Inbound Traffic Use Case and Topology with GWLB

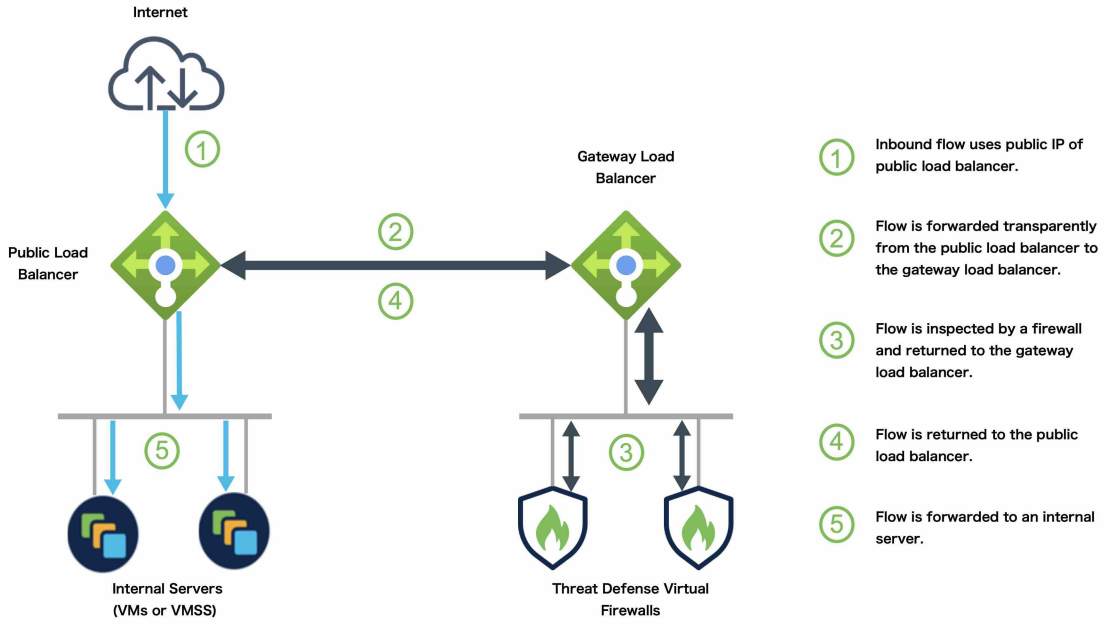

Figure 274: Outbound Traffic Use Case and Topology with GWLB

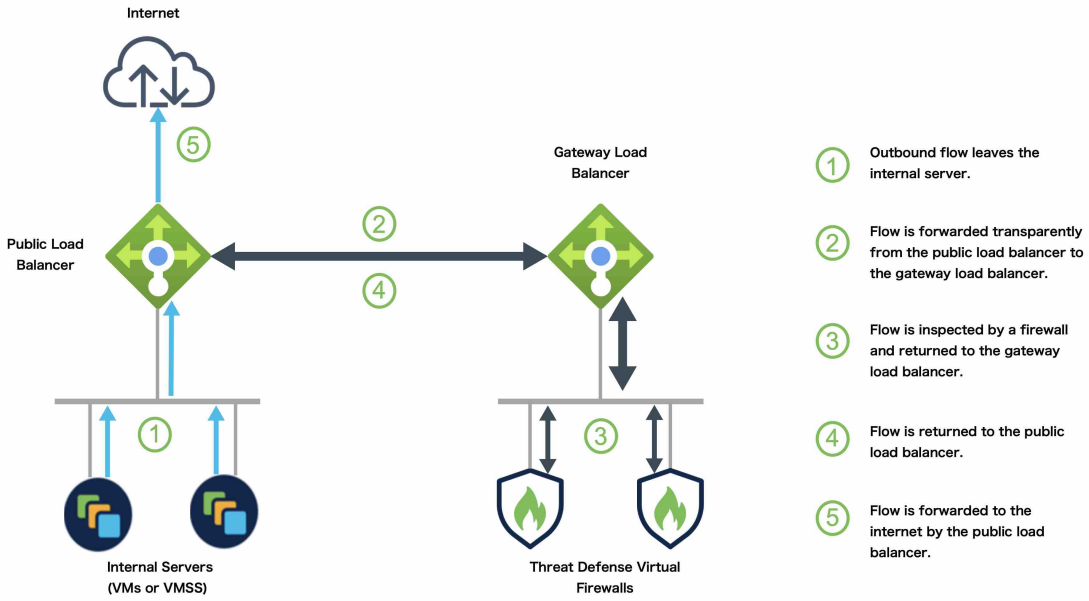

Azure Gateway Load Balancer and Paired Proxy

In an Azure service chain, Threat Defense Virtuals act as a transparent gateway that can intercept packets between the internet and the customer service. The Threat Defense Virtual defines an external interface and an internal interface on a single NIC by utilizing VXLAN segments in a paired proxy.

The following figure shows traffic forwarded to the Azure Gateway Load Balancer from the Public Gateway Load Balancer on the external VXLAN segment. The Gateway Load Balancer balances traffic among multiple Threat Defense Virtuals, which inspect the traffic before either dropping it or sending it back to the Gateway Load Balancer on the internal VXLAN segment. The Azure Gateway Load Balancer then sends the traffic back to the Public Gateway Load Balancer and to the destination.

Figure 275: Azure Gateway Load Balancer with Paired Proxy

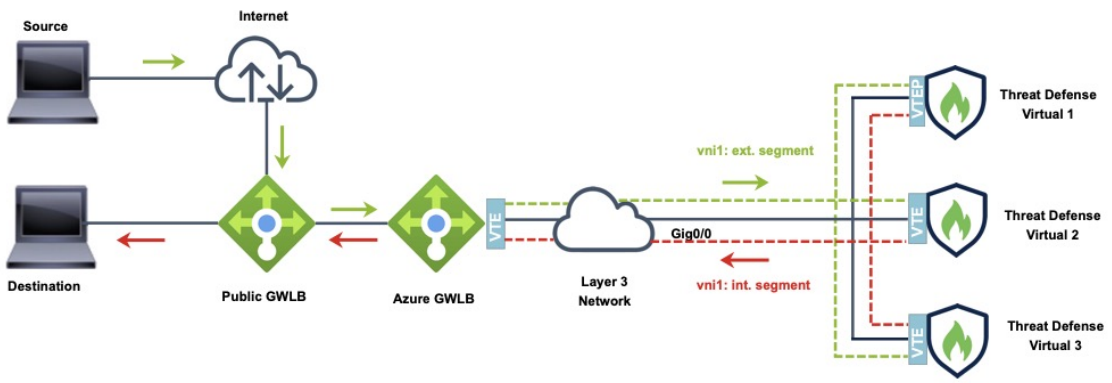

Traffic flow between GWLBe to GWLB (Geneve Single-Arm Proxy) in Azure

End-to-End Process for Deploying Threat Defense Virtual Cluster in Azure with GWLB

Template-based Deployment

The following flowchart illustrates the workflow for template-based deployment of the Threat Defense Virtual cluster in Azure with GWLB.

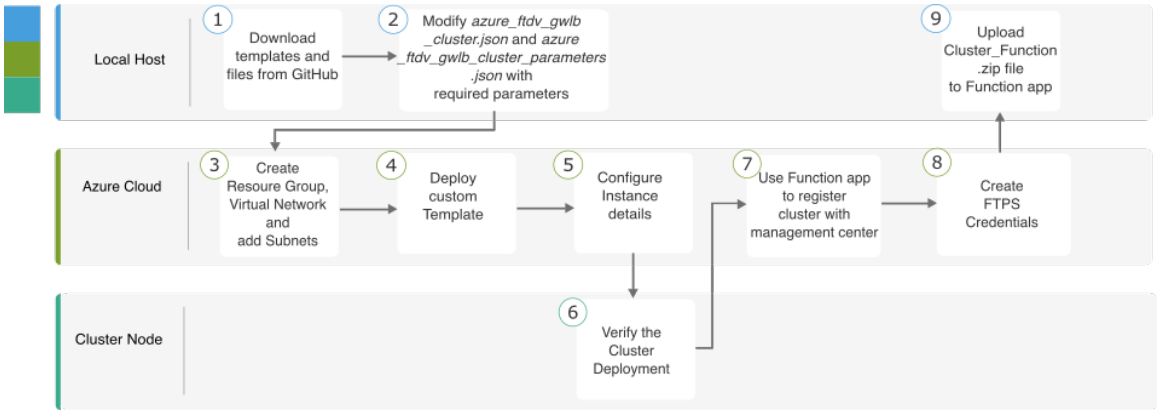

	Workspace	Steps
1	Local Host	Download templates and files from GitHub.
2	Local Host	Modify *azure_ftdv_gwlb_cluster.json* and *azure_ftdv_gwlb_cluster_parameters.json* with the required parameters.
3	Azure Cloud	Create the resource group, virtual network, and subnets.

	Workspace	Steps
4	Azure Cloud	Deploy custom template.
5	Azure Cloud	Configure instance details.
6	Cluster Node	Verify cluster deployment.
7	Azure Cloud	Use the Function app to register the cluster with the Management Center.
8	Azure Cloud	Create FTPS credentials.
9	Local Host	Upload *Cluster_Function.zip* file to the Function app.

Manual Deployment

The following flowchart illustrates the workflow of manual deployment of Threat Defense Virtual cluster in Azure with GWLB.

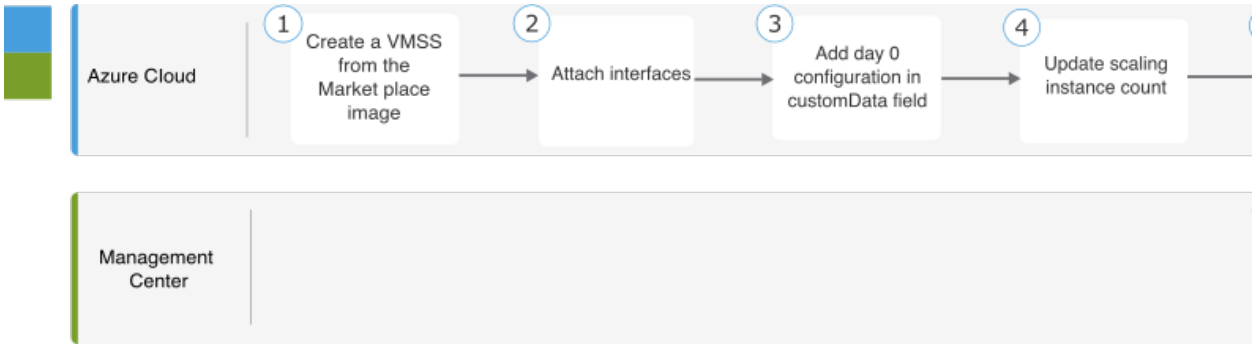

	Workspace	Steps
1	Local Host	Create a VMSS from the Marketplace image.
2	Local Host	Attach interfaces.
3	Local Host	Add day 0 configuration in the customData field.
4	Local Host	Update scaling instance count.
5	Local Host	Configure GWLB.
6	Management Center	Add control node.

Templates

The templates given below are available in GitHub. The parameter values are self-explanatory with the parameter names, and values, given in the template.

From Secure Firewall version 7.4.1, you can deploy the cluster without the diagnostic interface. To deploy the cluster with only the Outside, Inside, Management, and CCL interfaces, use the withoutDiagnostic templates - azure_withoutDiagnostic_ftdv_gwlb_cluster_parameters.json and azure_withoutDiagnostic_ftdv_gwlb_cluster.json files.

Templates to deploy **with Diagnostic** Interface:

- azure_ftdv_gwlb_cluster_parameters.json – Template to enter parameters for the Firewall Threat Defense Virtual cluster with GWLB.
- azure_ftdv_gwlb_cluster.json – Template to deploy Firewall Threat Defense Virtual cluster with GWLB.

Templates to deploy **without Diagnostic** Interface:

- azure_withoutDiagnostic_ftdv_gwlb_cluster_parameters.json – Template to enter parameters for the Firewall Threat Defense Virtual cluster with GWLB deployment without Diagnostic interface.
- azure_withoutDiagnostic_ftdv_gwlb_cluster.json – Template to deploy Firewall Threat Defense Virtual cluster with GWLB without Diagnostic interface.

Prerequisites

- To allow the cluster to auto-register to the management center, create a user with Network Admin & Maintenance User privileges on the management center. Users with these privileges can use REST API. See the Cisco Secure Firewall Management Center Administration Guide.
- Add an access policy in the management center that matches the name of the policy that you will specify during template deployment.
- Ensure that the Management Center Virtual is licensed appropriately.
- Perform the steps given below after the cluster is added to the Management Center Virtual:
  1. Configure platform settings with the health check port number in the Management Center. For more information on configuring this, see Platform Settings.
  2. Create a static route for data traffic. For more information on creating a static route, see Add a Static Route.

     Sample static route configuration:

```
Network: any-ipv4
Interface: vxlan_tunnel
Leaked from Virtual Router: Global
Gateway: vxlan_tunnel_gw
Tunneled: false
Metric: 2
```

Note *vxlan_tunnel_gw* is the data subnet's gateway IP address.

Deploy Cluster on Azure with GWLB Using an Azure Resource Manager Template

Deploy the Virtual Machine Scale Set for Azure GWLB using the customized Azure Resource Manager (ARM) template.

Procedure

Step 1 Prepare the template.

a) Clone the github repository to your local folder. See https://github.com/CiscoDevNet/cisco-ftdv/tree/master/cluster/azure.

b) Modify *azure_ftdv_gwlb_cluster.json* and *azure_ftdv_gwlb_cluster_parameters.json* with the required parameters.

OR

Modify withoutDiagnostic templates, *azure_withoutDiagnostic_ftdv_gwlb_cluster_parameters.json* and *azure_withoutDiagnostic_ftdv_gwlb_cluster.json*, with the required parameter for deploying cluster without the diagnostic interface.

Step 2 Log into the Azure Portal: https://portal.azure.com.

Step 3 Create a Resource Group.

a) In the **Basics** tab, choose the **Subscription** and **Resource Group** from the drop-down lists.

b) Choose the required **Region**.

Step 4 Create a virtual network with 4 subnets: Management, Diagnostic, Outside, and Cluster Control Link (CCL).

From Secure Firewall version 7.4.1, you can deploy the cluster without the diagnostic interface. To deploy the cluster with only the Outside, Inside, Management, and CCL interfaces, use the withoutDiagnostic templates - azure_withoutDiagnostic_ftdv_gwlb_cluster_parameters.json and azure_withoutDiagnostic_ftdv_gwlb_cluster.json files.

a) Create the virtual network.

1. In the **Basics** tab, choose the **Subscription** and **Resource Group** from the drop-down lists.
2. Choose the required **Region**. Click **Next: IP addresses**.

In the **IP Addresses** tab, click **Add subnet** and add the following subnets – Management, Diagnostic, Data, and Cluster Control Link.

If you are deploying the Firewall Threat Defense Virtual 7.4.1 cluster without a Diagnostic interface, then you must skip the Diagnostic subnet creation.

b) Add the subnets.

Step 5 Deploy the custom template.

a) Click **Create** > **Template deployment (deploy using custom templates)**.
b) Click **Build your own template in the editor**.
c) Click **Load File**, and upload **azure_ftdv_gwlb_cluster.json** or **azure_withoutDiagnostic_ftdv_gwlb_cluster.json**, if you have opted for without diagnostic interface deployment.
d) Click **Save**.

Step 6 Configure the Instance details.

a) Enter the required values and then click **Review + create**.
b) Click **Create** after the validation is passed.

Step 7 After the instance is running, verify the cluster deployment by logging into any one of the nodes and entering the **show cluster info** command.

Figure 276: show cluster info

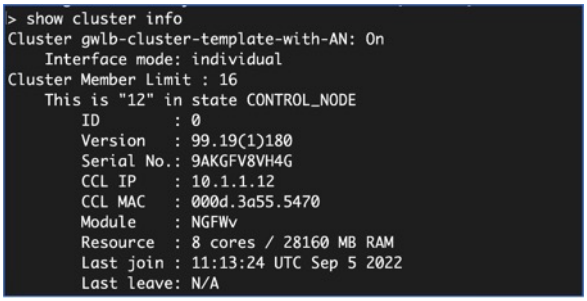

Step 8 In the Azure Portal, click the Function app to register the cluster with the Firewall Management Center.

Note
If you do not want to use the Function app, you can alternatively register the control node to the Firewall Management Center directly by using **Add** > **Device** (not **Add** > **Cluster**). The rest of the cluster nodes will register automatically.

Step 9 Create FTPS Credentials by clicking **Deployment Center** > **FTPS credentials** > **User scope** > **Configure Username and Password**, and then click **Save**.

Step 10 Upload the Cluster_Function.zip file to the Function app by executing the following **curl** command in the local terminal.

curl -X POST -u *username* **--data-binary @"Cluster_Function.zip" https://** *Function_App_Name***.scm.azurewebsites.net/api/zipdeploy**

Note
The **curl** command might take few minutes (~2 to 3 minutes) to complete command execution.

The function will be uploaded to the Function app. The function will start, and you can see the logs in the storage account's outqueue. The device registration with the Management Center will be initiated.

Figure 277: Functions

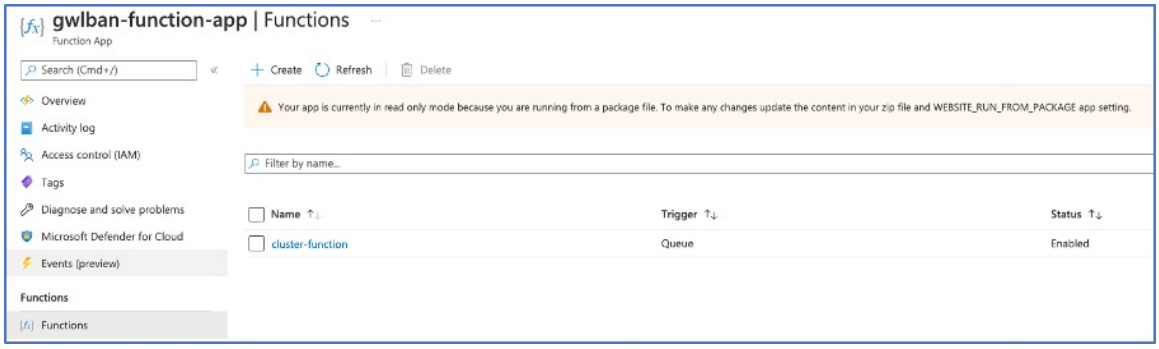

Figure 278: Queues

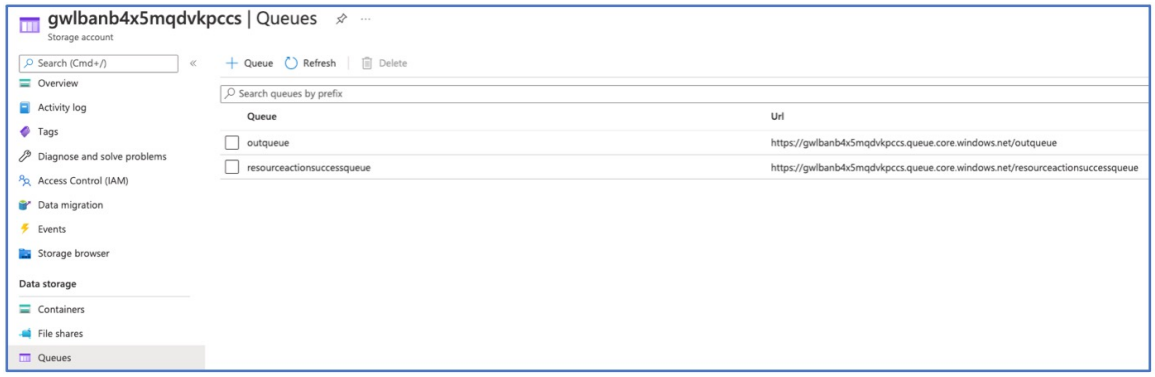

Figure 279: Outqueue

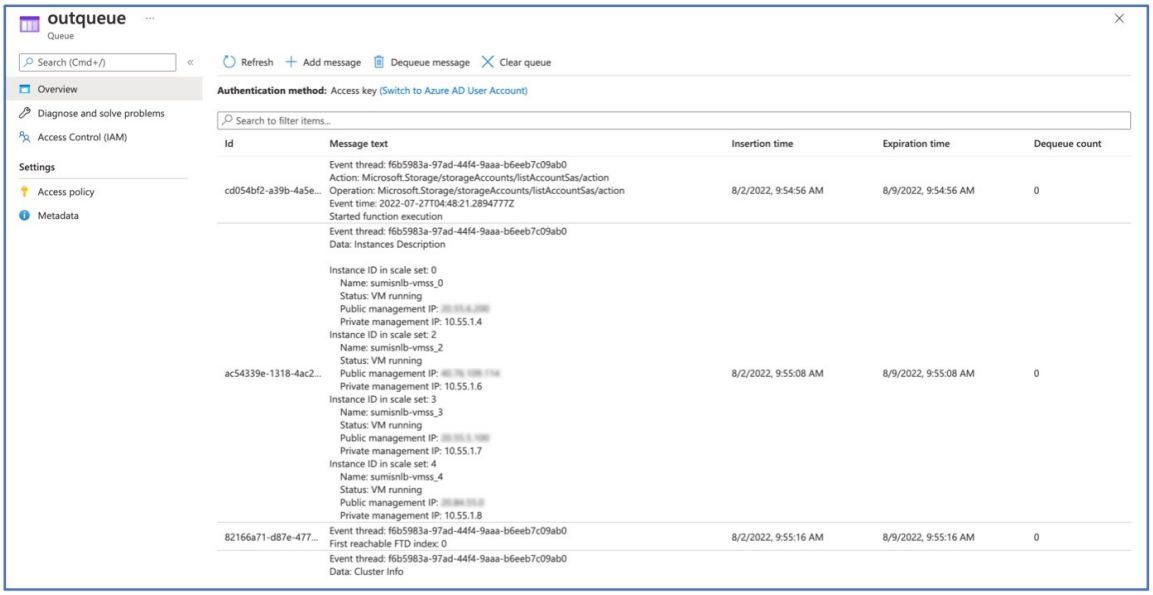

Sample Topology for NLB-based Cluster Deployment

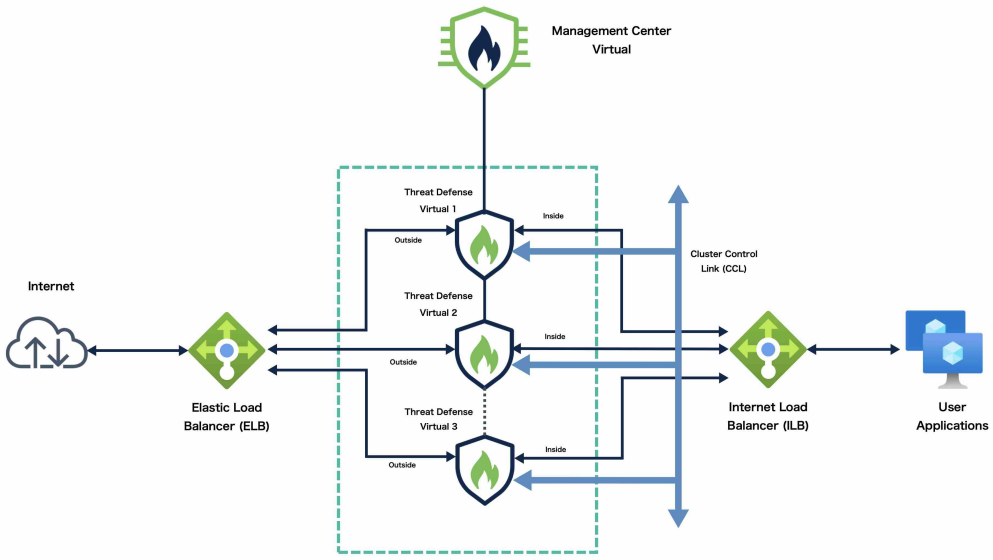

This topology depicts both inbound and outbound traffic flow. The Threat Defense Virtual cluster is sandwiched between the internal and external load balancers. A Management Center Virtual instance is used to manage the cluster.

Inbound traffic from the internet goes to the external load balancer which then transmits the traffic to the Threat Defense Virtual cluster. After the traffic has been inspected by a Threat Defense Virtual instance in the cluster, it is forwarded to the application VM.

Outbound traffic from the application VM is transmitted to the internal load balancer. Traffic is then forwarded to the Threat Defense Virtual cluster and then sent out to the internet.

End-to-End Process for Deploying Threat Defense Virtual Cluster in Azure with NLB

Template-based Deployment

The following flowchart illustrates the workflow of template-based deployment of Threat Defense Virtual cluster in Azure with NLB.

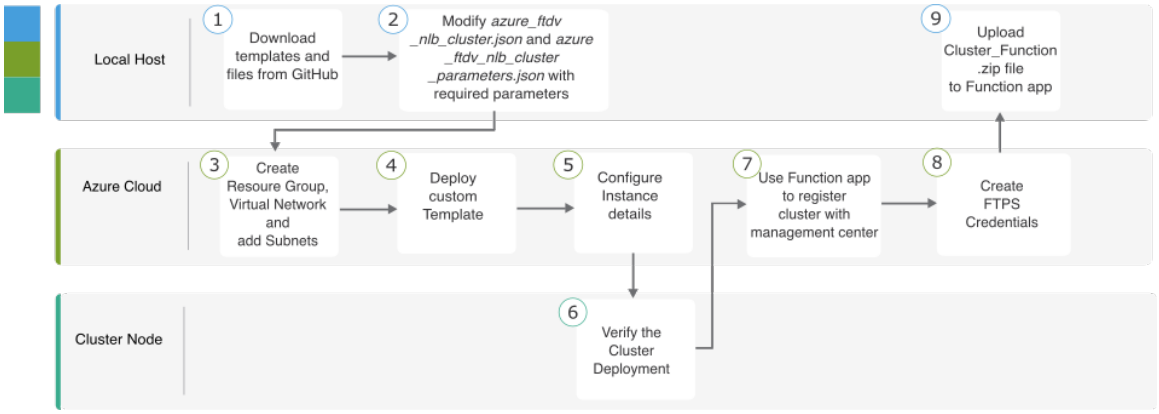

	Workspace	Steps
1	Local Host	Download templates and files from GitHub.
2	Local Host	Modify *azure_ftdv_nlb_cluster.json* and *azure_ftdv_nlb_cluster_parameters.json* with the required parameters.
3	Azure Cloud	Create the resource group, virtual network, and subnets.
4	Azure Cloud	Deploy custom template.
5	Azure Cloud	Configure instance details.
6	Cluster Node	Verify cluster deployment.
7	Azure Cloud	Use the Function app to register the cluster with the Management Center.
8	Azure Cloud	Create FTPS credentials.
9	Local Host	Upload *Cluster_Function.zip* file to the Function app.

Manual Deployment

The following flowchart illustrates the workflow of manual deployment of Threat Defense Virtual cluster in Azure with NLB.

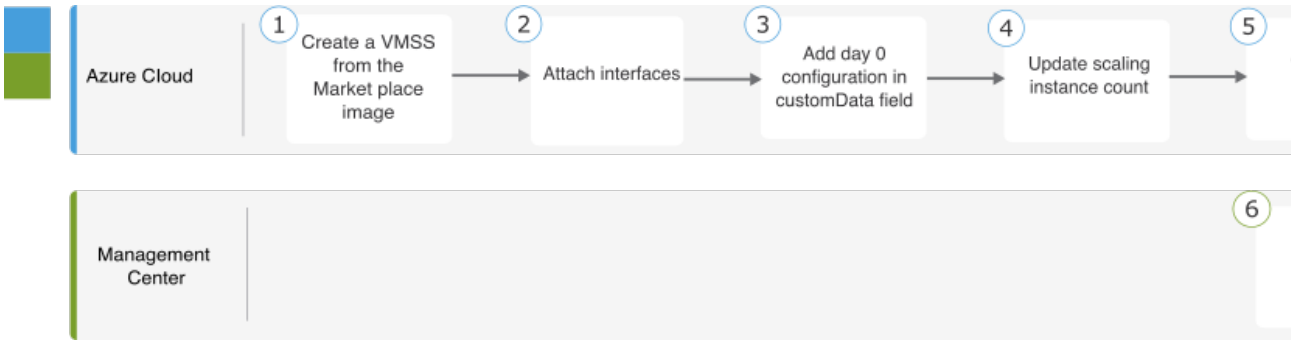

	Workspace	Steps
1	Local Host	Create a VMSS from the Marketplace image.
2	Local Host	Attach interfaces.
3	Local Host	Add day 0 configuration in the customData field.
4	Local Host	Update scaling instance count.
5	Local Host	Configure NLB.
6	Management Center	Add control node.

Templates

The templates given below are available in GitHub. The parameter values are self-explanatory with the parameter names, and values, given in the template.

From Secure Firewall version 7.4.1, you can deploy the cluster without the diagnostic interface. To deploy the cluster with only the Outside, Inside, Management, and CCL interfaces, use the withoutDiagnostic templates - azure_withoutDiagnostic_ftdv_nlb_cluster_parameters.json and azure_withoutDiagnostic_ftdv_nlb_cluster.json files.

Templates to deploy **with Diagnostic** Interface:

- azure_ftdv_nlb_cluster_parameters.json – Template to enter parameters for the Threat Defense Virtual cluster with NLB.
- azure_ftdv_nlb_cluster.json – Template to deploy Threat Defense Virtual cluster with NLB.

Templates to deploy **without Diagnostic** Interface:

- azure_withoutDiagnostic_ftdv_nlb_cluster_parameters.json – Template to enter parameters for the Firewall Threat Defense Virtual cluster with NLB deployment without Diagnostic interface.
- azure_withoutDiagnostic_ftdv_nlb_cluster.json – Template to deploy Firewall Threat Defense Virtual cluster with NLB without Diagnostic interface.

Prerequisites

- To allow the cluster to auto-register with the Management Center, create a user with Network Admin & Maintenance User privileges on the Management Center. Users with these privileges can use REST API. See the Cisco Secure Firewall Management Center Administration Guide.
- Add an access policy in the Management Center that matches the name of the policy that you will specify during template deployment.
- Ensure that the Management Center Virtual is licensed appropriately.
- After the cluster is added to the Management Center Virtual:

  1. Configure platform settings with the health check port number in the Management Center. For more information on configuring this, see Platform Settings.
  2. Create static routes for traffic from outside and inside interfaces. For more information on creating a static route, see Add a Static Route.

     Sample static route configuration for the outside interface:

     ```
     Network: any-ipv4
     Interface: outside
     Leaked from Virtual Router: Global
     Gateway: ftdv-cluster-outside
     Tunneled: false
     Metric: 10
     ```

     Note *ftdv-cluster-outside* is the outside subnet's gateway IP address.

     Sample static route configuration for the inside interface:

     ```
     Network: any-ipv4
     Interface: inside
     Leaked from Virtual Router: Global
     Gateway: ftdv-cluster-inside-gw
     Tunneled: false
     Metric: 11
     ```

     Note *ftdv-cluster-inside-gw* is the inside subnet's gateway IP address.

  3. Configure NAT rule for data traffic. For more information on configuring NAT rules, see Network Address Translation.

Deploy Cluster on Azure with NLB Using an Azure Resource Manager Template

Deploy the cluster for Azure NLB using the customized Azure Resource Manager (ARM) template.

Procedure

Step 1 Prepare the template.

a) Clone the github repository to your local folder. See https://github.com/CiscoDevNet/cisco-ftdv/tree/master/cluster/azure.

b) Modify *azure_ftdv_nlb_cluster.json* and *azure_ftdv_nlb_cluster_parameters.json* with the required parameters.

Modify withoutDiagnostic templates, *azure_withoutDiagnostic_ftdv_nlb_cluster_parameters.json* and *azure_withoutDiagnostic_ftdv_nlb_cluster.json*, with the required parameter for deploying cluster without the diagnostic interface.

Step 2 Log into the Azure Portal: https://portal.azure.com.

Step 3 Create a Resource Group.

a) In the **Basics** tab, choose the **Subscription** and **Resource Group** from the drop-down lists.

b) Choose the required **Region**.

Step 4 Create a virtual network with 5 subnets: Management, Diagnostic, Inside, Outside, and Cluster Control Link.

From Secure Firewall version 7.4.1, you can deploy the cluster without the diagnostic interface. To deploy the cluster with only the Outside, Inside, Management, and Cluster Control Link interfaces, use the withoutDiagnostic templates - azure_withoutDiagnostic_ftdv_nlb_cluster_parameters.json and azure_withoutDiagnostic_ftdv_nlb_cluster.json files.

a) Create the virtual network.

1. In the **Basics** tab, choose the **Subscription** and **Resource Group** from the drop-down lists.
2. b) Choose the required **Region**. Click **Next: IP addresses**.

b) Add the subnets.

In the **IP Addresses** tab, click **Add subnet** and add the following subnets – Management, Diagnostic, Inside, Outside, and Cluster Control Link.

If you are deploying the Firewall Threat Defense Virtual 7.4.1 cluster without a Diagnostic interface, then you must skip the Diagnostic subnet creation.

Step 5 Deploy the custom template.

a) Click **Create** > **Template deployment (deploy using custom templates)**.

b) Click **Build your own template in the editor**.

c) Click **Load File**, and upload **azure_ftdv_nlb_cluster.json** or **azure_withoutDiagnostic_ftdv_nlb_cluster.json**, if you have opted for without diagnostic interface deployment.

d) Click **Save**.

Step 6 Configure the instance details.

a) Enter the required values and then click **Review + create**.

Note
For the cluster control link starting and ending addresses, specify only as many addresses as you need (up to 16). A larger range can affect performance.

b) Click **Create** after the validation is passed.

Step 7 After the instance is running, verify the cluster deployment by logging into any one of the nodes and using the **show cluster info** command.

Figure 280: show cluster info

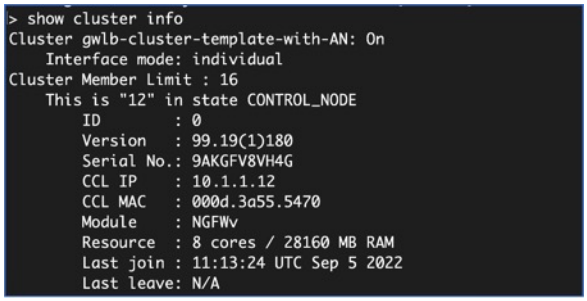

Step 8 In the Azure Portal, click the Function app to register the cluster to the Firewall Management Center.

Note
If you do not want to use the Function app, you can alternatively register the control node with the Management Center directly by using **Add** > **Device** (not **Add** > **Cluster**). The rest of the cluster nodes will register automatically.

Step 9 Create FTPS Credentials by clicking **Deployment Center** > **FTPS credentials** > **User scope** > **Configure Username and Password**, and then click **Save**.

Step 10 Upload the Cluster_Function.zip file to the Function app by executing the following **curl** command in the local terminal.

curl -X POST -u *username* **--data-binary @"Cluster_Function.zip" https://** *Function_App_Name***.scm.azurewebsites.net/api/zipdeploy**

Note
The **curl** command might take a few minutes (~2 to 3 minutes) to complete command execution.

The function will be uploaded to the Function app. The function will start, and you can see the logs in the storage account's outqueue. The device registration with the Management Center will be initiated.

Deploy the Cluster in Azure Manually

To deploy the cluster manually, prepare the day0 configuration, deploy each node, and then add the control node to the Firewall Management Center.

Create the Day0 Configuration for Azure

You can use either a fixed configuration or a customized configuration.

Create the Day0 Configuration With a Fixed Configuration for Azure

The fixed configuration will auto-generate the cluster bootstrap configuration.

```
"Cluster": {
"CclSubnetRange": "ip_address_start ip_address_end",
"ClusterGroupName": "cluster_name",
"HealthProbePort": "port_number",
"GatewayLoadBalancerIP": "ip_address",
"EncapsulationType": "vxlan",
"InternalPort": "internal_port_number",
"ExternalPort": "external_port_number",
"InternalSegId": "internal_segment_id",
"ExternalSegId": "external_segment_id"
}
```

Example

A sample day 0 configuration is given below.

```
"Cluster": {
"CclSubnetRange": "10.45.3.4 10.45.3.30",     //mandatory user input
"ClusterGroupName": "ngfwv-cluster",          //mandatory user input
"HealthProbePort": "7777",                    //mandatory user input
"GatewayLoadBalancerIP": "10.45.2.4",         //mandatory user input
"EncapsulationType": "vxlan",
"InternalPort": "2000",
"ExternalPort": "2001",
"InternalSegId": "800",
"ExternalSegId": "801"
}
```

Note If you are copying and pasting the configuration given above, ensure that you remove **//mandatory user input** from the configuration

For the Azure health check settings, be sure to specify the **HealthProbePort** you set here.

For the **CclSubnetRange** variable, specify a range of IP addresses starting from x.x.x.4. Ensure that you have at least 16 available IP addresses for clustering. Some examples of start and end IP addresses are given below.

Table 39: Examples of Start and End IP addresses

CIDR	Start IP Address	End IP Address
10.1.1.0/27	10.1.1.4	10.1.1.30
10.1.1.32/27	10.1.1.36	10.1.1.62
10.1.1.64/27	10.1.1.68	10.1.1.94
10.1.1.96/27	10.1.1.100	10.1.1.126
10.1.1.128/27	10.1.1.132	10.1.1.158
10.1.1.160/27	10.1.1.164	10.1.1.190
10.1.1.192/27	10.1.1.196	10.1.1.222
10.1.1.224/27	10.1.1.228	10.1.1.254

Create the Day0 Configuration With a Customized Configuration for Azure

You can enter the entire cluster bootstrap configuration using commands.

```
"Cluster": {
"CclSubnetRange": "ip_address_start ip_address_end",
"ClusterGroupName": "cluster_name",
"HealthProbePort": "port_number",
"GatewayLoadBalancerIP": "ip_address",
"EncapsulationType": "vxlan",
"InternalPort": "internal_port_number",
"ExternalPort": "external_port_number",
"InternalSegId": "internal_segment_id",
"ExternalSegId": "external_segment_id"
}
```

Deploy Cluster Nodes Manually - GWLB-based Deployment

Deploy the cluster nodes so they form a cluster.

Procedure

Step 1 Log into the Azure Portal: https://portal.azure.com

Step 2 Create a Resource Group.

a. In the **Basics** tab, choose the **Subscription** and **Resource Group** from the drop-down lists.

b. Choose the required **Region**.

Step 3 Create a Virtual Network with the necessary subnets: Management, Data and Cluster Control Link (CCL).

Note
Configure the CCL with the smallest subnet mask as required. Wider subnets can impact performance.

See the Azure document for creating the Virtual Network and subnet: https://learn.microsoft.com/en-us/azure/virtual-network/quickstart-create-virtual-network?tabs=portal

Step 4 Go to the Marketplace and search for **Cisco Secure Firewall Threat Defense Virtual – BYOL and PAYG** and click **Create**.

Step 5 Fill the required details and choose **Yes** for **Is this VM going to be part of Cluster?**

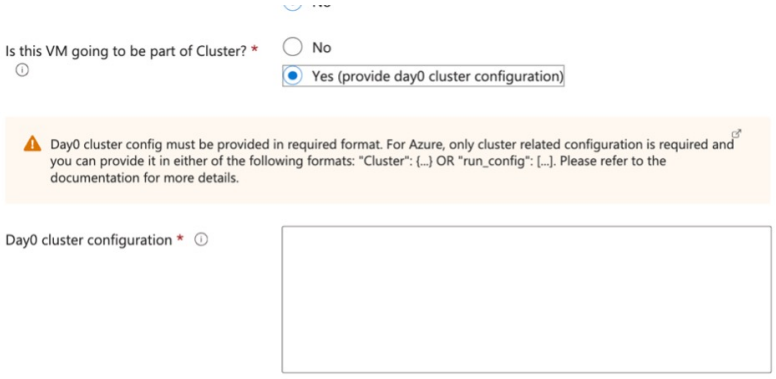

Paste the following cluster-related configuration in the text box.

```
"Cluster": {
"CclSubnetRange": "ip_address_start ip_address_end", //mandatory user input
"ClusterGroupName": "cluster_name", //mandatory user input
"HealthProbePort": "port_number", //mandatory user input
"GatewayLoadBalancerIP": "ip_address",  //mandatory user input
"EncapsulationType": "vxlan",
"InternalPort": "internal_port_number",
"ExternalPort": "external_port_number",
"InternalSegId": "internal_segment_id",
"ExternalSegId": "external_segment_id"
}
```

Step 6 Click **Next** and select the **Virtual Network & Subnets**.

Ensure **GigabitEthernet 0/1 subnet** is configured with the CCL subnet.

Step 7 Click **Review + create**. Wait until the Threat Defense Virtual deployment is completed.

Step 8 Connect to the Threat Defense Virtual device and execute **show cluster info** to confirm the cluster formation is successful.

```
> show cluster info
Cluster ngfwv-cluster: On
    Interface mode: individual
Cluster Member Limit : 16
    This is "4" in state CONTROL_NODE
        ID        : 0
        Version   : 9.23(1)
        Serial No.: 9AC1VMGJKAQ
        CCL IP    : 1.1.1.4
        CCL MAC   : 6045.bda8.e07b
        Module    : NGFWv
        Resource  : 4 cores / 14336 MB RAM
        Last join : 05:22:55 UTC Jul 14 2025
        Last leave: N/A
Other members in the cluster:
    There is no other unit in the cluster
>
```

Step 9 Configure the Azure Gateway Load Balancer. See Auto Scale with Azure Gateway Load Balancer Use Case for more information.

Step 10 Add the control node to the Firewall Management Center. See Add the Cluster to the Management Center (Manual Deployment), on page 662.

Deploy Cluster Nodes Manually - NLB-based Deployment

Deploy the cluster nodes so they form a cluster.

Procedure

Step 1 Log into the Azure Portal: https://portal.azure.com

Step 2 Create a Resource Group.

a. In the **Basics** tab, choose the **Subscription** and **Resource Group** from the drop-down lists.

b. Choose the required **Region**.

Step 3 Create a Virtual Network with the necessary subnets: Management, Inside, Outside and Cluster Control Link (CCL).

Note
Configure the CCL with the smallest subnet mask as required. Wider subnets can impact performance.

See the Azure document for creating the Virtual Network and subnet: https://learn.microsoft.com/en-us/azure/virtual-network/quickstart-create-virtual-network?tabs=portal

Step 4 Go to the Marketplace and search for **Cisco Secure Firewall Threat Defense Virtual – BYOL and PAYG** and click **Create**.

Step 5 Fill the required details and choose **Yes** for **Is this VM going to be part of Cluster?**

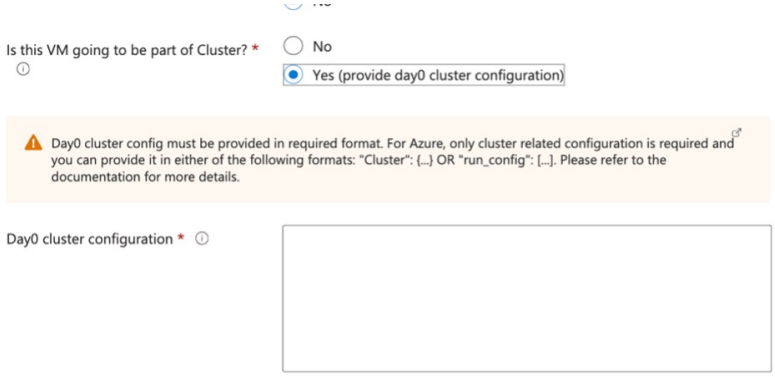

Paste the following cluster-related configuration in the text box.

```
"Cluster": {
"CclSubnetRange": "ip_address_start ip_address_end", //mandatory user input
"ClusterGroupName": "cluster_name" //mandatory user input
}
```

Step 6 Click **Next** and select the **Virtual Network & Subnets**.

If the diagnostic interface is enabled, user can attach maximum of four interfaces while deploying the Threat Defense Virtual VM. If the diagnostic is disabled, user can attach maximum of three interfaces while deploying the Threat Defense Virtual instance. So, the user must attach the extra interface for the cluster related communication after deploying the Threat Defense Virtual VM.

Step 7 Click **Review + create**. Wait until the Threat Defense Virtual deployment is completed.

Step 8 Power off the Threat Defense Virtual VM.

Step 9 Create a new interface using the CCL subnet and attach it to the Threat Defense Virtual VM. See the Azure document to create and attach new interface to the existing VM: https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-network-interface-vm

Step 10 Power On the Threat Defense Virtual VM.

Step 11 Connect to the Threat Defense Virtual device and execute **show cluster info** to confirm the cluster formation is successful.

```
> show cluster info
Cluster ngfwv-cluster: On
    Interface mode: individual
Cluster Member Limit : 16
    This is "4" in state CONTROL_NODE
        ID          : 0
        Version     : 9.23(1)
        Serial No.: 9AC1VMGJKAQ
        CCL IP      : 1.1.1.4
        CCL MAC     : 6045.bda8.e07b
        Module      : NGFWv
        Resource    : 4 cores / 14336 MB RAM
        Last join : 05:22:55 UTC Jul 14 2025
        Last leave: N/A
Other members in the cluster:
    There is no other unit in the cluster
>
```

Step 12 Add the control node to the Management Center. See Add the Cluster to the Management Center (Manual Deployment), on page 662.

Troubleshooting Cluster Deployment in Azure

- Issue: No traffic flow

  Troubleshooting:

  - Check if the health probe status of the Threat Defense Virtual instances deployed with a GWLB is healthy.
  - If the Threat Defense Virtual instance's health probe status is unhealthy-
    - Check if the static route is configured in the Management Center Virtual.
    - Check if the default gateway is the data subnet's gateway IP.
    - Check if the Threat Defense Virtual instance is receiving health probe traffic.
    - Check if the access list configured in the Management Center Virtual allows health probe traffic.

- Issue: Cluster is not formed

  Troubleshooting:

  - Check the IP address of the nve-only cluster interface. Ensure that you can ping the nve-only cluster interface of other nodes.
  - Check the IP address of the nve-only cluster interfaces are part of the object group.
  - Ensure that the NVE interface is configured with the object group .
  - Ensure that the cluster interface in the cluster group has the right VNI interface. This VNI interface has the NVE with the corresponding object group.
  - Ensure that the nodes are pingable from each other. Since each node has its own cluster interface IP, these should be pingable from each other.
  - Check if the CCL Subnet's Start and End Address mentioned during template deployment is correct. The start address should begin with the first available IP address in the subnet. For example, if the subnet is 192.168.1.0/24. The start address should be 192.168.1.4 (the three IP addresses at the start are reserved by Azure).
  - Check if the Management Center Virtual has a valid license.

- Issue: Role-related error while deploying resources again in the same resource group.

  Troubleshooting: Remove the roles given below by using the following commands on the terminal.

  Error message:

```
"error": {
"code": "RoleAssignmentUpdateNotPermitted",
"message": "Tenant ID, application ID, principal ID, and scope are not allowed to be
updated.”}
```

  - **az role assignment delete --resource-group** <*Resource Group Name*> **--role "Storage Queue Data Contributor"**
  - **az role assignment delete --resource-group** <*Resource Group Name*> **--role "Contributor"**

Firewall Threat Defense Virtual Clustering Autoscale Solution on Azure

A typical cluster deployment in an Azure region includes a defined number of Firewall Threat Defense Virtual instances (nodes). When the Azure region traffic varies, without dynamic scaling (autoscale) of the nodes, resource utilization in such cluster arrangement may underutilise the resources or cause latency. Cisco offers an autoscale solution for Firewall Threat Defense Virtual clustering in Version 7.7 and later that supports dynamic scaling of nodes in the Azure region. It allows you to scale-in or scale-out nodes from the cluster based on the network traffic. It uses logic based on the resource utilization statistics from Azure VMSS metrics such as CPU and memory metrics to dynamically add or remove a node from a cluster.

The Firewall Threat Defense Virtual clustering with Autoscale solution in Azure supports both Network Load Balancer (NLB or Sandwich topology) and Gateway Load Balancer (GWLB). See Sample Topologies, on page 630

Cisco provides separate Azure Resource Manager (ARM) templates for deploying Firewall Threat Defense Virtual cluster with autoscale in Azure using NLB and GWLB, as well as infrastructure and configuration templates for deploying the Azure services such as Function App and Logic App.

Sample Topologies

Firewall Threat Defense Virtual Clustering with Autoscale in Azure using Sandwich Topology (Network Load Balancer)

The Firewall Threat Defense Virtual clustering with autoscale in Azure using sandwich topology (NLB) use case is an automated horizontal scaling solution that positions the Firewall Threat Defense Virtual scale set sandwiched between an Azure Internal load balancer (ILB) and an Azure External load balancer (ELB).

In this topology, the Firewall Threat Defense Virtual uses only *four* interfaces: management, inside, outside, and CCL subnets.

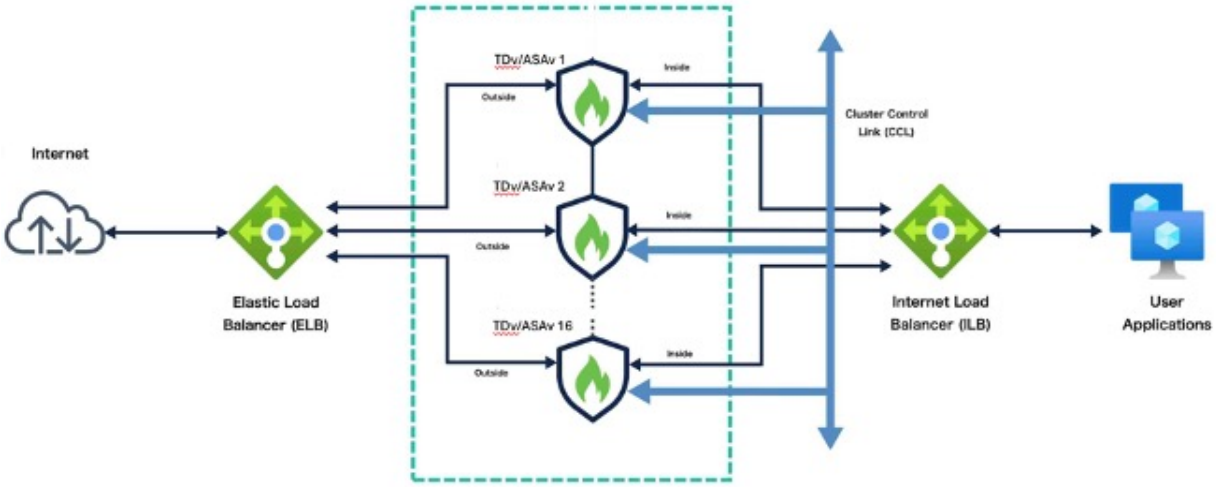

Firewall Threat Defense Virtual Clustering with Autoscale in Azure using Sandwich Topology (NLB)

The following describes high-level flow on how a Firewall Threat Defense Virtual cluster with autoscale in Azure using NLB functions:

- The ELB distributes traffic from the internet to the Firewall Threat Defense Virtual instances in the scale set, and then the firewall forwards traffic to the application.
- The ILB distributes outbound internet traffic from an application to Firewall Threat Defense Virtual instances in the scale set and then the firewall forwards traffic to the internet.
- A network packet will never pass through both (Internal and External) load balancers in a single connection.
- The number of Firewall Threat Defense Virtual instances in the scale set will be scaled and configured automatically based on load conditions.

Firewall Threat Defense Virtual Clustering with Autoscale in Azure using Gateway Load Balancer

The integration of the Azure Gateway Load Balancer (GWLB) and Firewall Threat Defense Virtual cluster using autoscale solution simplifies deployment, management, and scaling of instances in the cluster setup. The Azure Gateway Load Balancer (GWLB) ensures that internet traffic to and from an Azure VM, such as an application server, is inspected by secure firewall without requiring any routing changes. This integration also reduces operational complexity and provides a single entry and exit point for traffic at the firewall. The applications and infrastructure can maintain visibility of source IP address, which is critical in some environments.

The Firewall Threat Defense Virtual uses only *three* interfaces: management, data, and CCL interface in this use case.

Note

- Network Address Translation (NAT) is not required if you are deploying the Azure GWLB.
- Only IPv4 is supported.

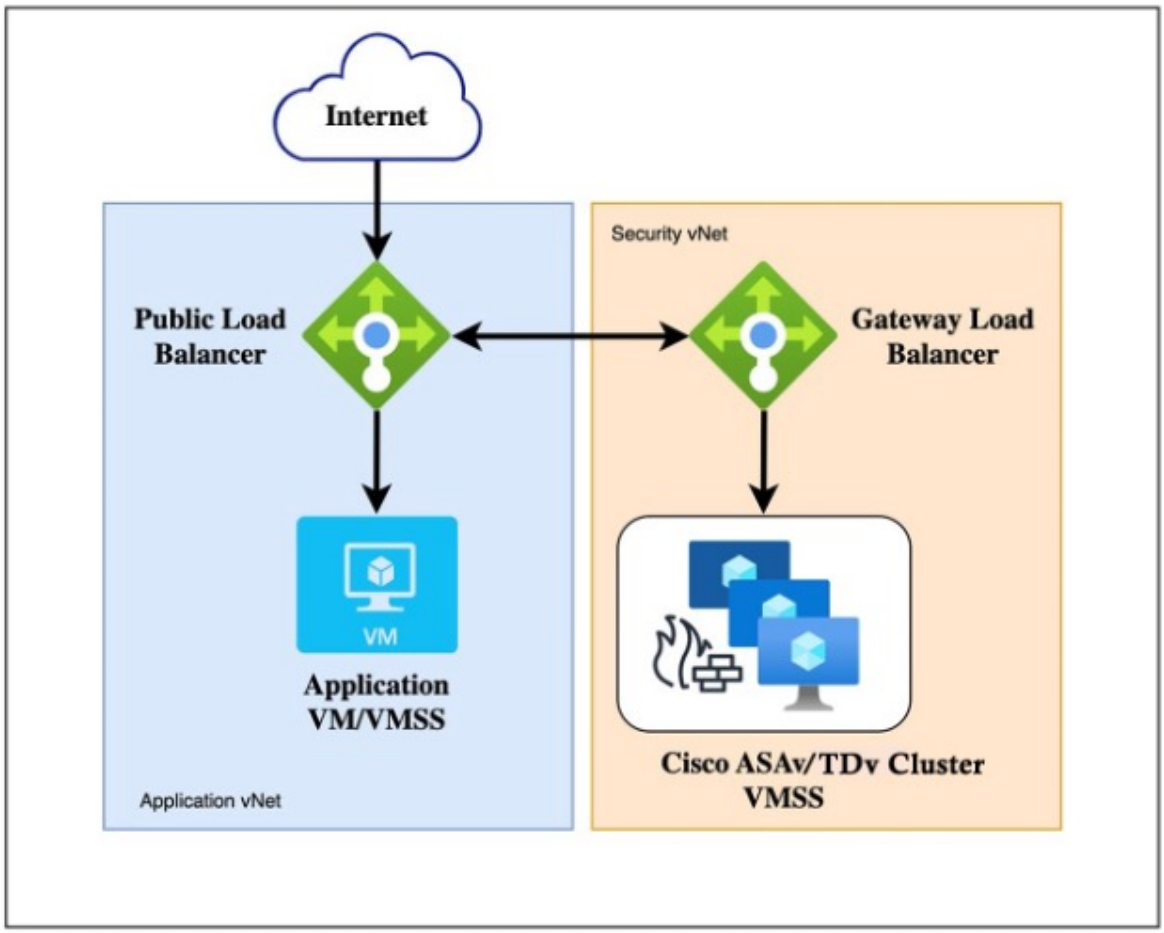

The following describes high-level flow on how a Firewall Threat Defense Virtualcluster with autoscale in Azure using GWLB functions:

- Inbound traffic from the internet goes to the GWLB endpoint, which then transmits the traffic to the GWLB.
- The traffic is then routed to the Firewall Threat Defense Virtual cluster.
- After the traffic is inspected by the Firewall Threat Defense Virtual instance in the cluster, it is forwarded to the application Application VM.

Prerequisites

- Ensure that you have Owner role in the Azure subscription.
- Create the Azure Resource Group. Ensure that the Azure Virtual Network along with the necessary subnets are created.
  - Interfaces for NLB-based cluster : Management, Diagnostic, Inside, Outside, CCL and the function app.
  - Interfaces for GWLB-based cluster : Management, Diagnostic, Data, CCL and the function app.
- On the Management Center:
  - Ensure that Management Center Virtual is licensed correctly.
  - Create the access control policy.
  - Create the Security Zone (SZ) object for the interfaces. For NLB based cluster, create the SZ for inside and outside interfaces. For GWLB-based cluster, create the SZ for the data interface.
  - Create a separate user name and password for the azure function to add the Threat Defense Virtual instances to the Management Center Virtual and configure the instances.
- Install the Azure CLI on your local system.
- Download the Azure Clustering Autoscale repository from GitHub to your local computer and run the command **python3 make.py build** to create the Azure functions zip file.

Autoscale Logic for Firewall Threat Defense Virtual Clustering in Azure

Scaling Policy

In a cluster with autoscale, the scaling of nodes is determined based on the following policies:

- Scaling policy 1 - When one cluster node exceeds the resource utilization limits.
- Scaling policy 2 - Overall average resource utilization of all the nodes.

Scale-out

Scale-out is a process of adding a new node to the cluster when the traffic load threshold exceeds the configured CPU or memory limit on any one of the cluster's node.

The following is the process of adding a new node to the cluster during scale-out:

1. A new Firewall Threat Defense Virtual instance is launched.

2. Appropriate configuration is applied to a Firewall Threat Defense Virtual.
3. Appropriate licenses are applied.
4. A new Firewall Threat Defense Virtual instance is added to the cluster.

If the configuration of the new Firewall Threat Defense Virtual instance fails (low probability) during the scale-out process, the failing instance is terminated, and a new instance is launched and configured.

Scale-in

Scale-in is the process of removing a node from a cluster when the configured scale-in threshold and total number of cluster instances exceed the minimum cluster size.

The following is the process of terminating a node in the cluster during scale-in:

1. The Firewall Threat Defense Virtual instance with the least CPU or memory usage is identified using VMSS metrics.
2. If there is more than one instance with the same least utilization, then the instance with the higher VM index in VMSS is chosen for scale-in.
3. Any new connections to this instance are disabled by appropriate configuration and policies.
4. The instance is de-registered from smart licensing (applicable for BYOL).
5. The instance is terminated.

Azure Functions (Function App)

The Function application helps to enable the Firewall Threat Defense Virtual cluster and register it with the management center. The Function application also help you select a hosting plan for Firewall Threat Defense Virtual clustering with autoscale deployment.

The following two types of hosting plans are offered:

- **Consumption**
  - This is the default hosting plan for Firewall Threat Defense Virtual clustering with autoscale.
  - This plan allows the Function app to connect to the Firewall Threat Defense Virtual instances by opening the SSH port to the Azure data center IP addresses of the region.
- **Premium**
  - You can select this hosting plan for the Function app during deployment.
  - This plan supports adding a Network Address Translation (NAT) gateway to the Function app to control the outbound IP address of the Function app. This plan allows SSH access to Firewall Threat Defense Virtual instances only from a fixed IP address of the NAT gateway thereby offering enhanced security.

For more information about overview on auto scale solution components, see Auto Scale Solution Components in *Cisco Secure Firewall Threat Defense Virtual Getting Started Guide*.

Deployment and Infrastructure Templates on GitHub

Cisco provides Azure Resource Manager (ARM) templates and scripts for deploying an auto-scaling group of Firewall Threat Defense Virtual cluster using several Azure services, including Function App, Logic App, auto-scaling groups and so on.

The autoscale solution for Firewall Threat Defense Virtual cluster is an ARM template-based deployment that provides:

- Completely automated Firewall Threat Defense Virtual instance registration and de-registration with the management center using the Function App.
- NAT policy, access control policy, and routes automatically applied to the scaled-out threat defense virtual instances.
- Support for GWLB and NLB load balancers.
- Works only with the management center; the device manager is not supported.

Firewall Threat Defense Virtual Clustering with Autoscale Solution Templates

Azure Resource Manager (ARM) templates

Two sets of templates are provided for autoscale solutions based on the (NLB or GWLB) load balancer you are using in Azure for the cluster.

The following templates are available on GitHub:

- Autoscale solution template for Firewall Threat Defense Virtual clustering using NLB: `azure_ftdv_nlb_cluster.json.json` available in the folder `arm-templates`.
- Autoscale solution template for Firewall Threat Defense Virtual clustering using GWLB: `azure_ftdv_gwlb_cluster.json` available in the folder `arm-templates`.

Setting up Azure Infrastructure and Configuration

- Function app to enable cluster on Firewall Threat Defense Virtual instances: `cluster_functions.zip`.
- Logic App code for the Firewall Threat Defense Virtual deployment, scale-in and scale-out workflow: `logic_app.txt`.

Input Parameters

The following table defines the template parameters and provides an example. Once you decide on these values, you can use these parameters to create the Firewall Threat Defense Virtual when you deploy the Azure Resource Manager (ARM) template into your Azure subscription.In the clustering with autoscale soultion with GWLB for Azure, networking infrastructure is also created due to which additional input parameters have to be configured in the template. The parameter descriptions are self-explanatory.

Table 40: Template Parameters

Parameter Name	Allowed Values/Type	Description	Resource Creation Type
resourceNamePrefix	String* (3-10 characters)	All the resources are created with name containing this prefix. Note: Use only lowercase letters. Example: ftdv	New
virtualNetworkRg	String	The virtual network resource group name. Example: cisco-virtualnet-rg	Existing
virtualNetworkName	String	The virtual network name (already created). Example: cisco-virtualnet	Existing
virtualNetworkCidr	CIDR format x.x.x.x/y	CIDR of Virtual Network (already created)	Existing
mgmtSubnet	String	The management subnet name (already created). Example: cisco-mgmt-subnet	Existing
dataSubnet	String	The data subnet name (already created) Example: cisco-data-subnet	
cclSubnet	String	The cluster control link subnet name. Example: cisco-ccl-subnet	
cclSubnetStartAddr	String	The starting range of CCL subnet IP address. Example: 3.4.5.6	
cclSubnetEndAddr	String	The ending range of CCL subnet IP address. Example: 5.6.7.8	
gwlbIP	String	GWLB is created in existing data subnet. Example: 10.0.2.4	
dataNetworkGatewayIp	String	The gateway IP address of the data subnet. Example: 10.0.2.7	

Parameter Name	Allowed Values/Type	Description	Resource Creation Type
outsideSecurityZoneName	String	The security zone object Name created in the management center Example: outside-sz	
TDvmManagementUserName	String	TDv management administrator username. You are not allowed provide 'admin' as the username.	
diagSubnet	String	The diagnostic subnet name (already created). Example: cisco-diag-subnet	Existing
insideSubnet	String	The inside Subnet name (already created). Example: cisco-inside-subnet	Existing
internalLbIp	String	The internal load balancer IP address for the inside subnet (already created). Example: 1.2.3.4	Existing
insideNetworkGatewayIp	String	The inside subnet gateway IP address (already created).	Existing
outsideSubnet	String	The outside subnet name (already created). Example: cisco-outside-subnet	Existing
outsideNetworkGatewayIp	String	The outside subnet gateway IP (already created).	Existing
deviceGroupName	String	Device group in Firewall Management Center (already created)	Existing
insideZoneName	String	Inside Zone name in the Firewall Management Center (already created)	Existing
outsideZoneName	String	Outside Zone name in the Firewall Management Center (already created)	Existing
softwareVersion	String	The Firewall Threat Defense Virtual Version (selected from drop-down list during deployment).	Existing

Parameter Name	Allowed Values/Type	Description	Resource Creation Type
vmSize	String	Size of Firewall Threat Defense Virtual instance (selected from drop-down list during deployment).	N/A
ftdLicensingSku	String	Firewall Threat Defense Virtual Licensing Mode (PAYG/BYOL) Note: PAYG is supported in Version 6.5+.	N/A
licenseCapability	Comma-separated string	BASE, MALWARE, URLFilter, THREAT	N/A
tdVmManagementUserName	String*	The Firewall Threat Defense Virtual VM management administrator user name. This **cannot** be 'admin'. See Azure for VM administrator user name guidelines.	New
tdVmManagementUserPassword	String*	Password for the Firewall Threat Defense Virtual VM management administrator user. Passwords must be 12 to 72 characters long, and must have: lowercase, uppercase, numbers, and special characters; and must have no more than 2 repeating characters. **Note** There is no compliance check for this in the template.	New
ftdAdminUserPassword	String	Firewall Threat Defense Virtual Admin user password. **Note** The criteria mentioned for the**TDvmManagementUserPassword** parameter is applicable to this parameter also.	
fmcIpAddress	String x.x.x.x	The public IP address of the Firewall Management Center (already created)	Existing

Parameter Name	Allowed Values/Type	Description	Resource Creation Type
fmcUserName	String	Firewall Management Center user name, with administrative privileges (already created)	Existing
fmcPassword	String	Firewall Management Center password for above Firewall Management Center user name (already created)	Existing
policyName	String	Security Policy created in the Firewall Management Center (already created)	Existing
clusterGroupName	String	The name of the cluster group to be used while registering the threat defense device to the management center. Example: tdv-cluster	
healthCheckPortNumber	String	The health check port number used while creating the health probe in the Gateway Load balancer. Example: 8080	
functionHostingPlan	String	Function deployment hosting plan (consumption uses the consumption hosting plan, premium: uses the premium hosting plan). Default: consumption	
functionAppSubnet	String	The function app subnet name (already created). Example: tdv-fapp-subnet	
functionAppSubnetCIDR	String	The CIDR of the function app subnet (already created). Example: 10.0.4.0/24	
scalingMetricsList	String	The metrics used in determining the scaling the scaling decision. Allowed: CPU & MEMORY	

Parameter Name	Allowed Values/Type	Description	Resource Creation Type
scalingPolicy	POLICY-1 / POLICY-2	**POLICY-1**: Scale-Out will be triggered when the average load of any Firewall Threat Defense Virtual goes beyond the Scale Out threshold for the configured duration. **POLICY-2**: Scale-Out will be triggered when average load of all the Firewall Threat Defense Virtual devices in the VMSS goes beyond the Scale Out threshold for the configured duration. In both cases Scale-In logic remains the same: Scale-In will be triggered when average load of all the Firewall Threat Defense Virtual devices comes below the Scale In threshold for the configured duration.	N/A
scalingMetricsList	String	Metrics used in making the scaling decision. Allowed: CPU, MEMORY Default: CPU	N/A
cpuScaleInThreshold	String	The scale-in threshold in percentage for CPU metrics. Default: 10 When the Firewall Threat Defense Virtual metric goes below this value the scale-in will be triggered. See Autoscale Logic for Firewall Threat Defense Virtual Clustering in Azure, on page 632 .	N/A

Parameter Name	Allowed Values/Type	Description	Resource Creation Type
cpuScaleOutThreshold	String	The Scale-out threshold in percentage for CPU metrics. Default: 80 When the Firewall Threat Defense Virtual metric goes above this value, the Scale-Out will be triggered. The 'cpuScaleOutThreshold' should always be **greater** than the 'cpuScaleInThreshold'. See Autoscale Logic for Firewall Threat Defense Virtual Clustering in Azure, on page 632.	N/A
memoryScaleInThreshold	String	The Scale-In threshold in percent for memory metrics. Default: 0 When the Firewall Threat Defense Virtual metric goes below this value the Scale-In will be triggered. See Autoscale Logic for Firewall Threat Defense Virtual Clustering in Azure, on page 632.	N/A
memoryScaleOutThreshold	String	The Scale-Out threshold in percent for memory metrics. Default: 0 When the Firewall Threat Defense Virtual metric goes above this value, the Scale-Out will be triggered. The 'memoryScaleOutThreshold' should always be **greater** than the 'memoryScaleInThreshold'. See Autoscale Logic for Firewall Threat Defense Virtual Clustering in Azure, on page 632.	N/A
minFtdCount	Integer	The minimum Firewall Threat Defense Virtual instances available in the scale set at any given time. Example: 2	N/A

Parameter Name	Allowed Values/Type	Description	Resource Creation Type
maxFtdCount	Integer	The maximum Firewall Threat Defense Virtual instances allowed in the Scale set. Example: 10 **Note** This number is restricted by the Firewall Management Center capacity. The Auto Scale logic will not check the range of this variable, hence fill this carefully.	N/A
metricsAverageDuration	Integer	Select from the drop-down. This number represents the time (in minutes) over which the metrics are averaged out. If the value of this variable is 5 (i.e. 5min), when the Auto Scale Manager is scheduled it will check the past 5 minutes average of metrics and based on this it will make a scaling decision. **Note** Only numbers 1, 5, 15, and 30 are valid due to Azure limitations.	N/A

Parameter Name	Allowed Values/Type	Description	Resource Creation Type
initDeploymentMode	BULK / STEP	Primarily applicable for the first deployment, or when the Scale Set does not contain any Firewall Threat Defense Virtual instances. BULK: The Auto Scale Manager will try to deploy 'minFtdCount' number of Firewall Threat Defense Virtual instances in parallel at one time. **Note** The launch is in parallel, but registering with the Firewall Management Center is sequential due to Firewall Management Center limitations. STEP: The Auto Scale Manager will deploy the 'minFtdCount' number of Firewall Threat Defense Virtual devices one by one at each scheduled interval. **Note** The STEP option will take a long time for the ‘minFtdCount’ number of instances to be launched and configured with the Firewall Management Center and become operational, but useful in debugging. The BULK option takes same amount of time to launch all ‘minFtdCount’ number of Firewall Threat Defense Virtual as one Firewall Threat Defense Virtual launch takes (because it runs in parallel), but the Firewall Management Center registration is sequential. The total time to deploy ‘minFtdCount’ number of Firewall Threat Defense Virtual = (time to launch One Firewall Threat Defense Virtual + time to register/configure one Firewall Threat Defense Virtual * minFtdCount).	

Parameter Name	Allowed Values/Type	Description	Resource Creation Type
*Azure has restrictions on the naming convention for new resources. Review the limitations or simply use all lowercase. **Do not use spaces or any other special characters**.			

Firewall Threat Defense Virtual Cluster with Autoscale Deployment Process and Resources

Firewall Threat Defense Virtual cluster with autoscale deployment process on Azure involves the following:

- Deploy the ARM template.
- Build and deploy the clustering function.
- Update and enable the Logic application.

Azure Resource Manager Template Deployment Resources

The following resources are created within a resource group when you deploy Firewall Threat Defense Virtual cluster with autoscale in Azure using the ARM template for **Sandwich Topology (NLB)** - `azure_ftdv_nlb_cluster_autoscale.json`

- Virtual Machine Scale Set (VMSS)
- External Load Balancer
- Internal Load Balancer
- Azure Function App
- Logic App
- Security groups (For Data and Management interfaces)

The following resources are created within a resource group when you deploy Firewall Threat Defense Virtual cluster with autoscale in Azure using the ARM template for GWLB - `azure_ftdv_gwlb_cluster_autoscale.json`

- Virtual Machine (VM) or Virtual Machine Scale Set (VMSS)
- Gateway Load Balancer (GWLB)
- Azure Function App
- Logic App
- Networking Infrastructure
- Security Groups and other miscellaneous components needed for deployment.

Deploy the Firewall Threat Defense Virtual Cluster with Autoscale Solution

Deploy the Threat Defense Virtual clustering with autoscale solution on Azure using the ARM template. Based on the topology, Sandwich (NLB) or GWLB use case, you are required to download and configure the appropriate ARM template for deploying the Firewall Threat Defense Virtual clustering with autoscale solution on Azure.

Before you begin

Download the Deployment Package from GitHub

The Firewall Threat Defense Virtual clustering autoscale with NLB solution for Azure is an Azure Resource Manager (ARM) template-based deployment which makes use of the serverless infrastructure provided by Azure (Logic App, Azure Functions, Load Balancers, Virtual Machine Scale Set, and so on).

The Firewall Threat Defense Virtual clustering autoscale with GWLB solution for Azure is an ARM template-based deployment that creates the GWLB, networking infrastructure, threat defense virtual auto scaling group, serverless components, and other required resources.

The deployment procedures for both solutions are similar.

Download the files required to launch the Firewall Threat Defense Virtual clustering with autoscale solution for Azure.

Deployment scripts and templates for your version are available in the GitHub repository.

Procedure

Step 1 Log in to the Microsoft Azure portal (https://portal.azure.com) using your Microsoft account username and password.

Step 2 Click **Resource groups** from the menu of services to access the **Resource Groups** blade. You will see all the resource groups in your subscription listed in the blade. Create a new resource group or select an existing, empty resource group. For example, **threat defense virtual_AutoScale**.

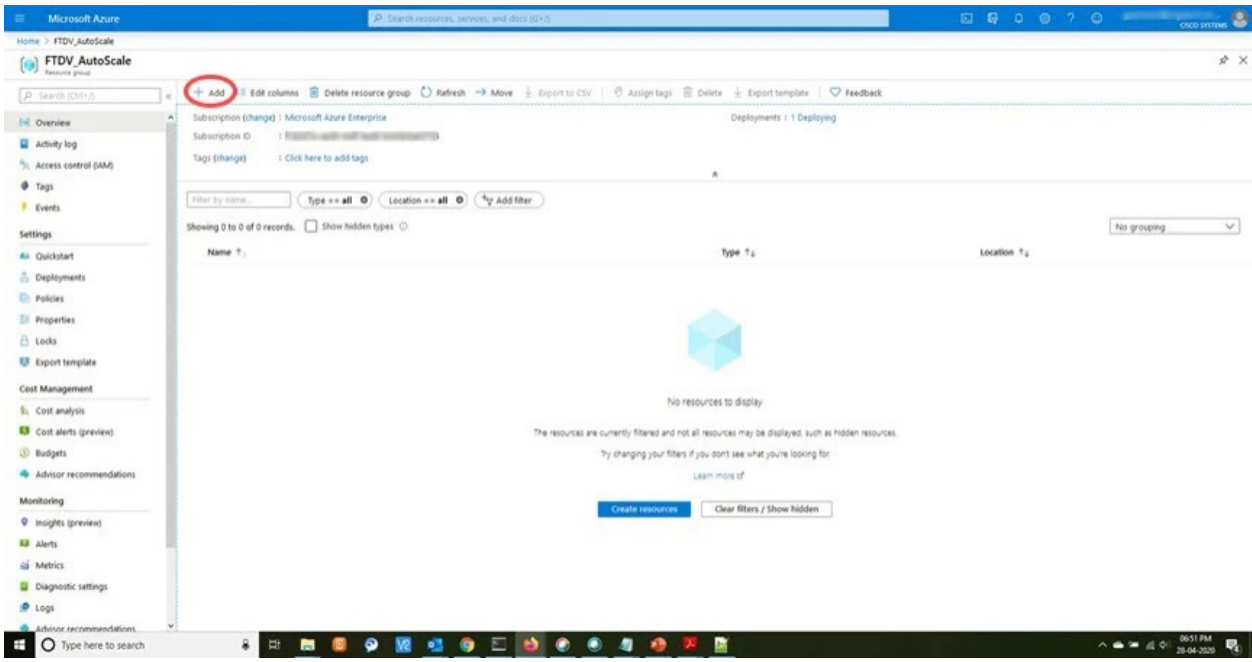

Step 3 Click **Create a resource (+)** to create a new resource for template deployment. The **Create Resource Group** blade appears.

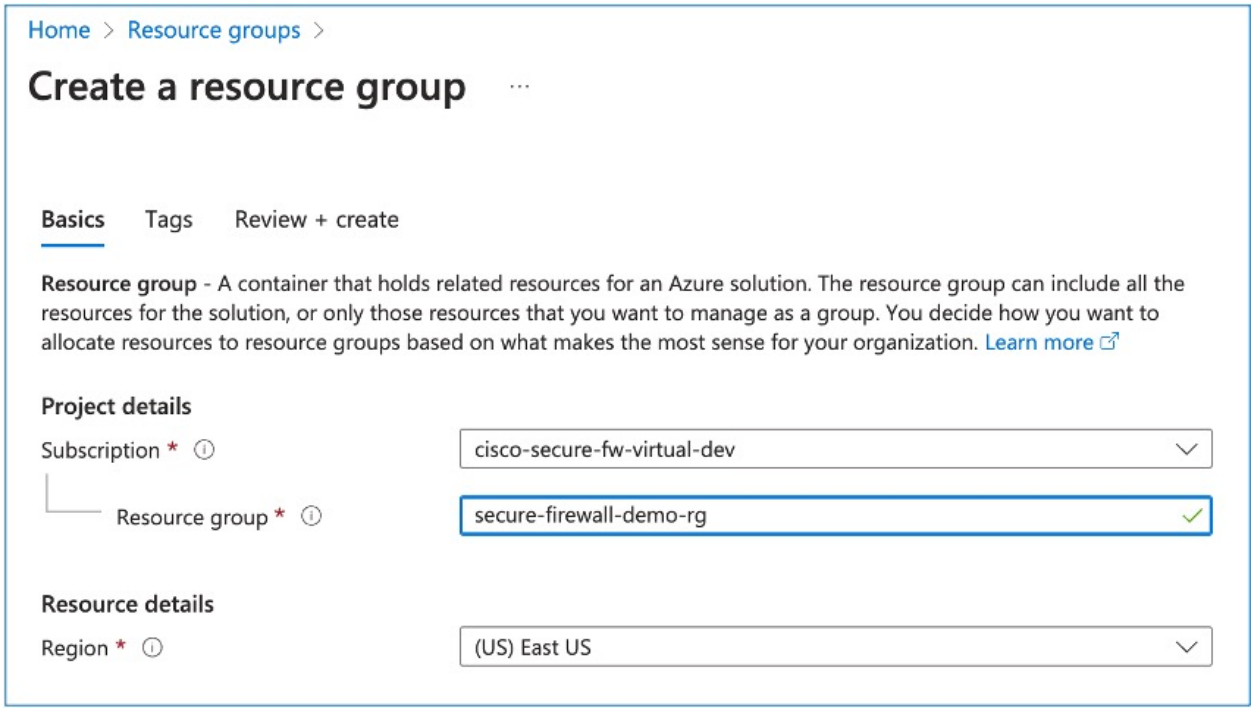

Step 4 4. Click **Virtual Network** from the menu of services to access the Virtual network blade. Create a virtual network with subnets.

- For GWLB deployment, create virtual network with management, data, CCL subnets, and the function app.

- For NLB deployment, create virtual network with management, inside, outside, CCL subnets and the function app.

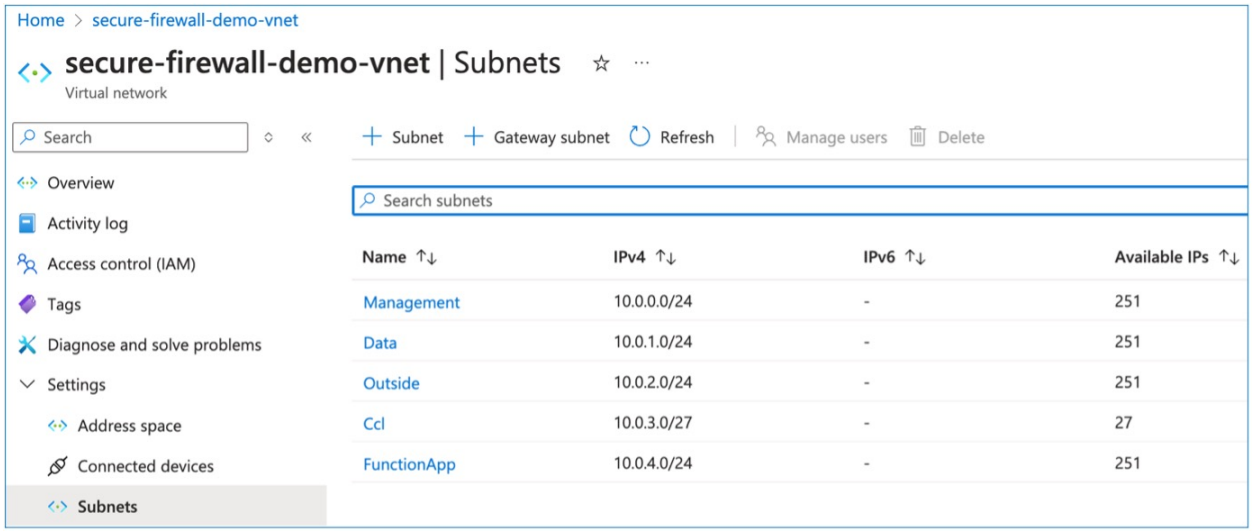

Step 5 In **Search the Marketplace**, type **`Template deployment`** (deploy using custom templates), and then press **Enter**.

Step 6 Click **Create**. There are several options for creating a template. Choose **Build your own template in editor**.

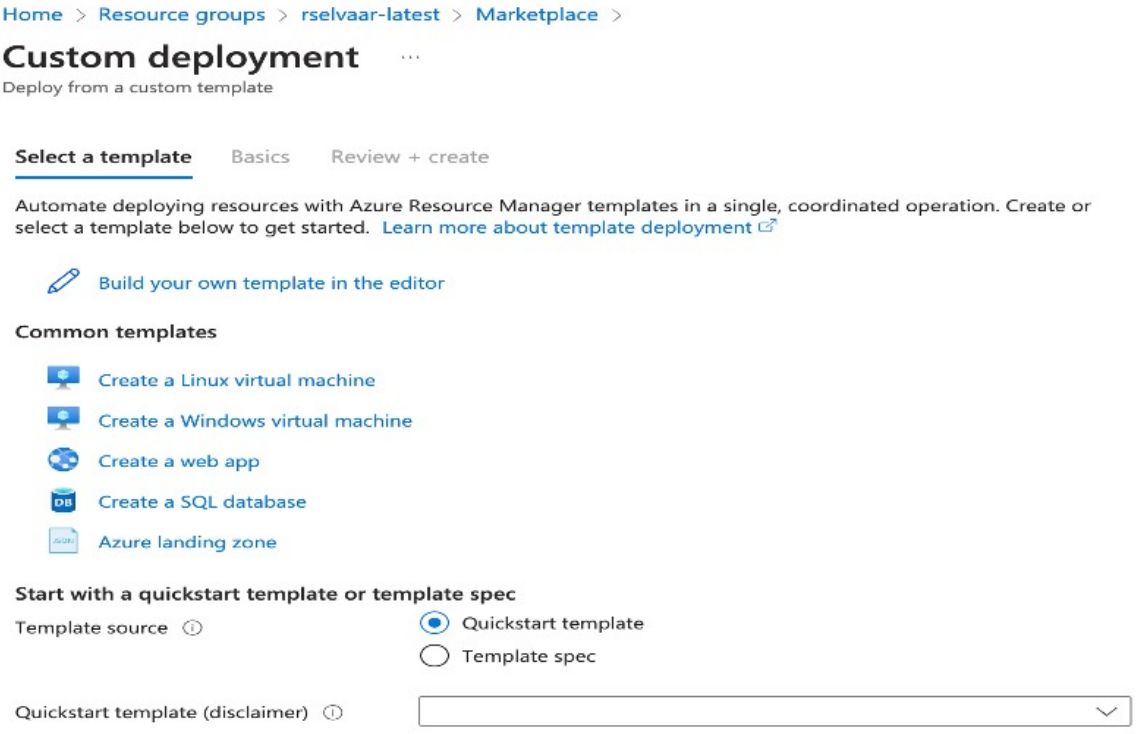

Step 7 In the **Edit template** window, delete all the default content and copy the contents from the updated `azure_ftdv_gwlb_cluster_custom_image.json` or

`azure_ftdv_nlb_cluster_custom_image.json` (depending on the type of autoscale solution you are deploying on Azure) and click **Save**. Or Click **Load file** to browse and upload this file from your computer.

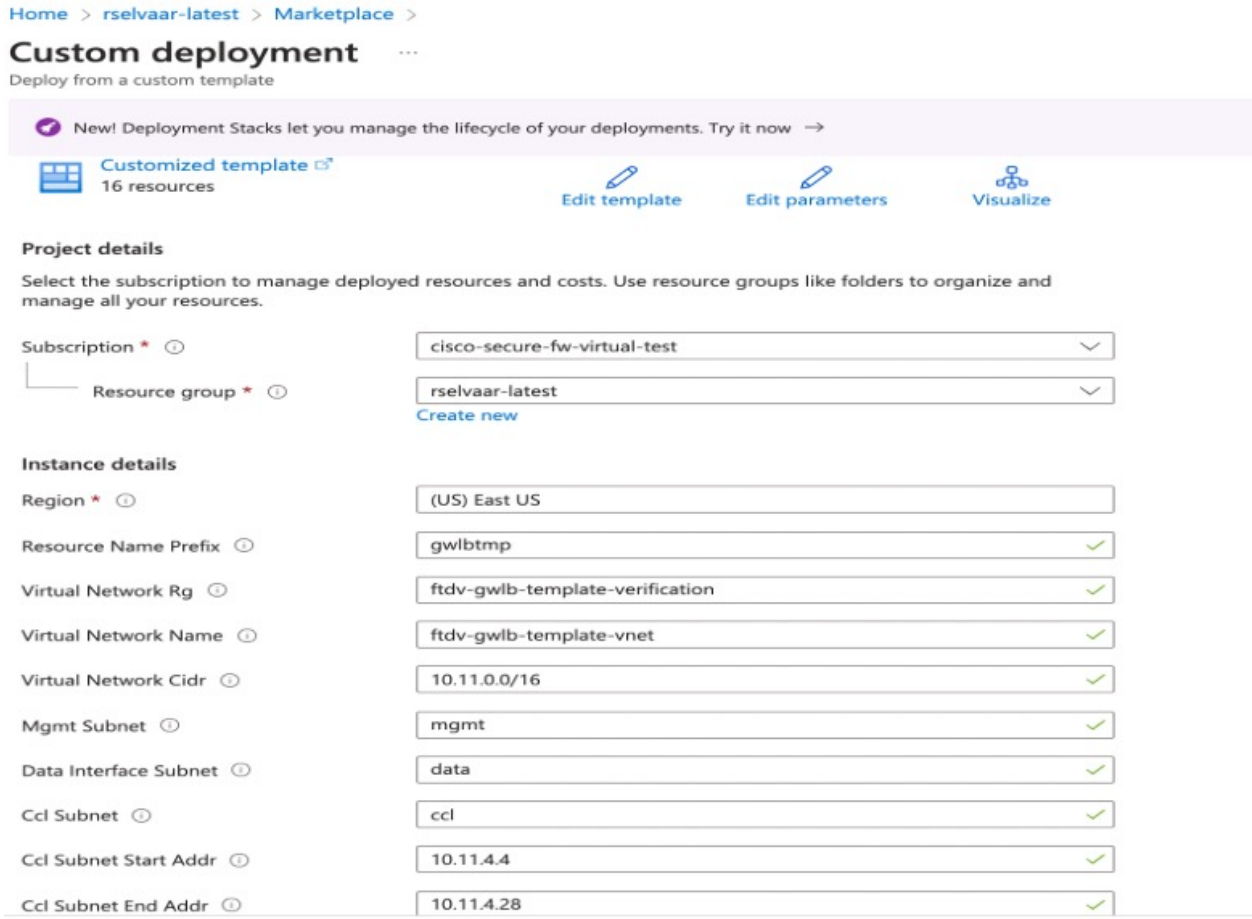

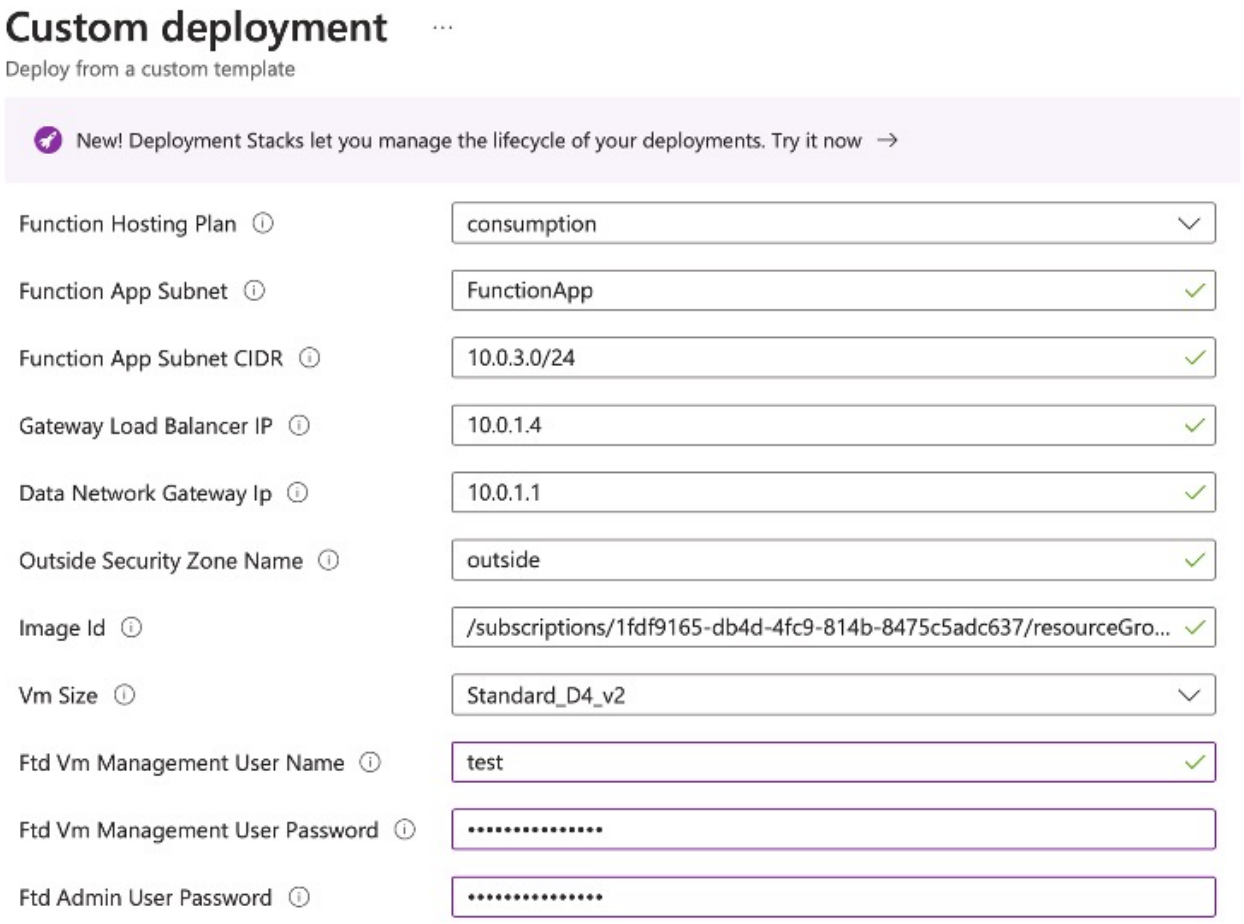
Custom deployment
Deploy from a custom template
New! Deployment Stacks let you manage the lifecycle of your deployments. Try it now →
Function Hosting Plan
consumption
Function App Subnet
FunctionApp
Function App Subnet CIDR
10.0.3.0/24
Gateway Load Balancer IP
10.0.1.4
Data Network Gateway Ip
10.0.1.1
Outside Security Zone Name
outside
Image Id
/subscriptions/1fdf9165-db4d-4fc9-814b-8475c5adc637/resourceGro...
Vm Size
Standard_D4_v2
Ftd Vm Management User Name
test
Ftd Vm Management User Password
Ftd Admin User Password

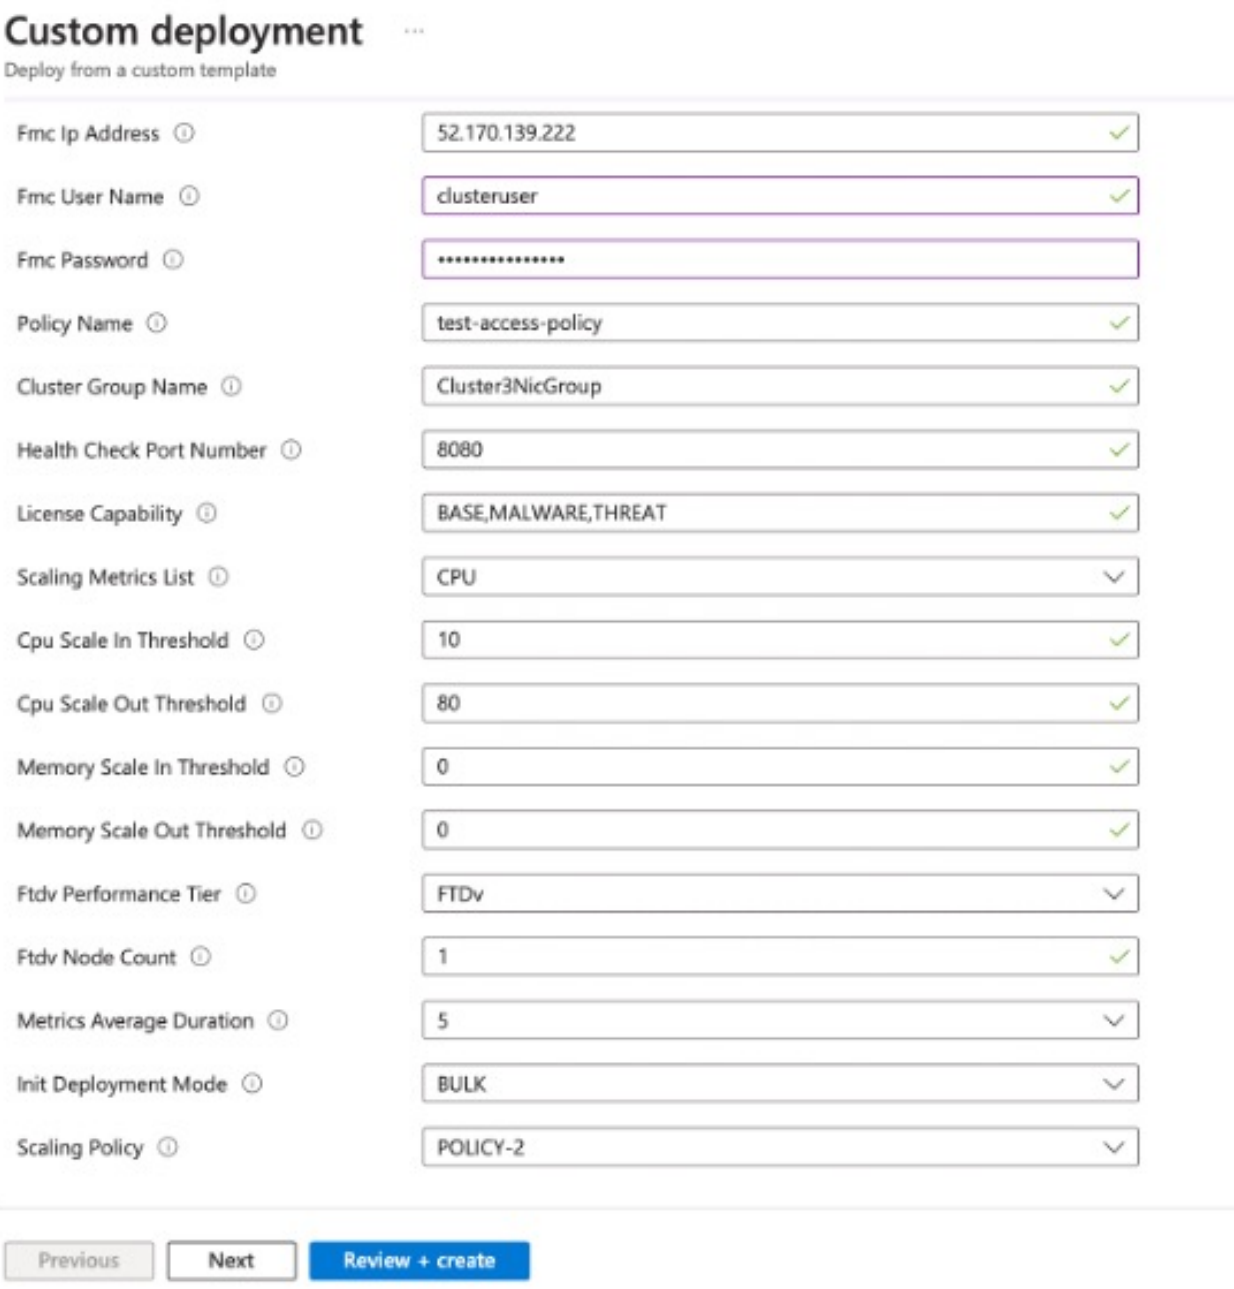

Step 8 In the parameter field sections, fill all the parameters. Refer to Input Parameters for details about each parameter, then click **Review+Create**.

Step 9 When a template deployment is successful, it creates all the required resources for the threat defense virtual auto scale for Azure solution. See the resources in the following figure. The **Type** column describes each resource, including the Logic App, VMSS, Load Balancers, Public IP address, etc.

What to do next

Deploy Azure Functions App, on page 650.

Deploy Azure Functions App

When you deploy the ARM template, Azure creates the function app with the name **<resourceNamePrefix>-function-app**.

Procedure

Step 1 Go to the function app you created when you deployed the ARM template and perform the following:

Run the following command from your local computer to deploy the cluster autoscale Azure Functions to the Function app.

```
az functionapp deployment source config-zip -g <Resource Group Name>
-n <Function App Name> --src  <cluster_functions.zip> --build-remote true
```

Step 2 After the deployment of the Azure Functions, you can view the uploaded Functions in the overview section of the function application.

Update the Azure Logic App

The Logic App acts as the orchestrator for the Autoscale functionality. The ARM template creates a skeleton Logic App, which you then need to update manually to provide the information necessary to function as the auto scale orchestrator.

Procedure

Step 1 From the repository, retrieve the file *LogicApp.txt* to the local system and edit as shown below.

Important
Read and understand all of these steps before proceeding.

These manual steps are not automated in the ARM template so that only the Logic App can be upgraded independently later in time.

a) Required: Find and replace all the occurrences of "SUBSCRIPTION_ID" with your subscription ID information.
b) Required: Find and replace all the occurrences of "RG_NAME" with your resource group name.
c) Required: Find and replace all of the occurrences of "FUNCTIONAPPNAME" to your function app name.

The following example shows a few of these lines in the *LogicApp.txt* file:

```
"AutoScaleManager": {
    "inputs": {
        "function": {
            "id":
```

```
"/subscriptions/SUBSCRIPTION_ID/resourceGroups/RG_NAME/providers/Microsoft.Web/sites/FUNCTIONAPPNAME/functions/AutoScaleManager"

          }
.
.
                    },
                    "Deploy_Changes_to_FTD": {
                        "inputs": {
                            "body": "@body('AutoScaleManager')",
                            "function": {
                                "id":
"/subscriptions/SUBSCRIPTION_ID/resourceGroups/RG_NAME/providers/Microsoft.Web/sites/FUNCTIONAPPNAME/functions/DeployConfiguration"

                            }
.
.
                    "DeviceDeRegister": {
                        "inputs": {
                            "body": "@body('AutoScaleManager')",
                            "function": {
                                "id":
"/subscriptions/SUBSCRIPTION_ID/resourceGroups/RG_NAME/providers/Microsoft.Web/sites/FUNCTIONAPPNAME/functions/DeviceDeRegister"

                            }
                        },
                        "runAfter": {
                            "Delay_For_connection_Draining": [
```

d) (Optional) Edit the trigger interval, or leave the default value (5). This is the time interval at which the Autoscale functionality is periodically triggered. The following example shows these lines in the *LogicApp.txt* file:

```
"triggers": {
    "Recurrence": {
        "conditions": [],
        "inputs": {},
        "recurrence": {
            "frequency": "Minute",
            "interval": 5
        },
```

e) (Optional) Edit the time to drain, or leave the default value (5). This is the time interval to drain existing connections from the Firewall Threat Defense Virtual before deleting the device during the Scale-In operation. The following example shows these lines in the *LogicApp.txt* file:

```
"actions": {
     "Branch_based_on_Scale-In_or_Scale-Out_condition": {
         "actions": {
             "Delay_For_connection_Draining": {
                 "inputs": {
                     "interval": {
                         "count": 5,
                         "unit": "Minute"
                     }
```

f) (Optional) Edit the cool down time, or leave the default value (10). This is the time to perform NO ACTION after the Scale-Out is complete. The following example shows these lines in the *LogicApp.txt* file:

```
"actions": {
    "Branch_based_on_Scale-Out_or_Invalid_condition": {
        "actions": {
            "Cooldown_time": {
                "inputs": {
                    "interval": {
                        "count": 10,
                        "unit": "Second"
                    }
```

Note
These steps can also be done from the Azure portal. Consult the Azure documentation for more information.

Step 2 Go to the **Logic App code view**, delete the default contents and paste the contents from the edited *LogicApp.txt* file, and click **Save**.

Figure 281: Logic App Code View

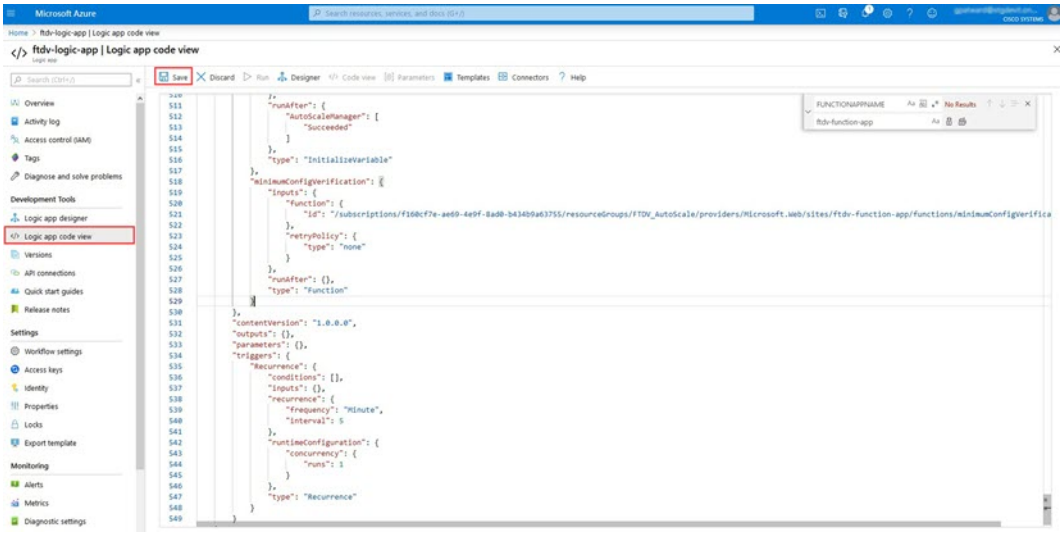

Step 3 When you save the Logic App, it is in a 'Disabled' state. Click **Enable** when you want to start the Auto Scale Manager.

Figure 282: Enable Logic App

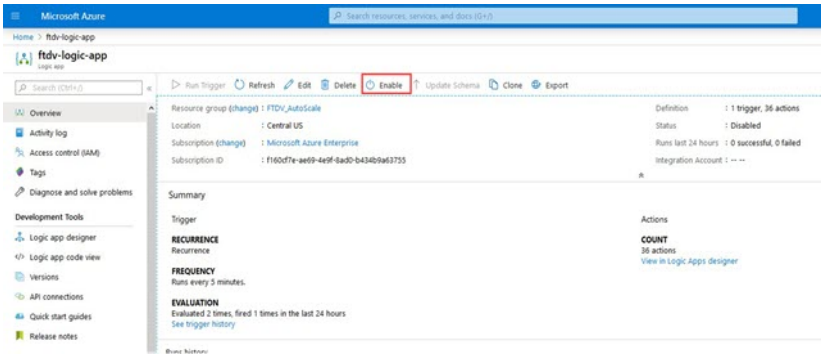

Step 4 Once enabled, the tasks start running. Click the 'Running' status to see the activity.

Figure 283: Logic App Running Status

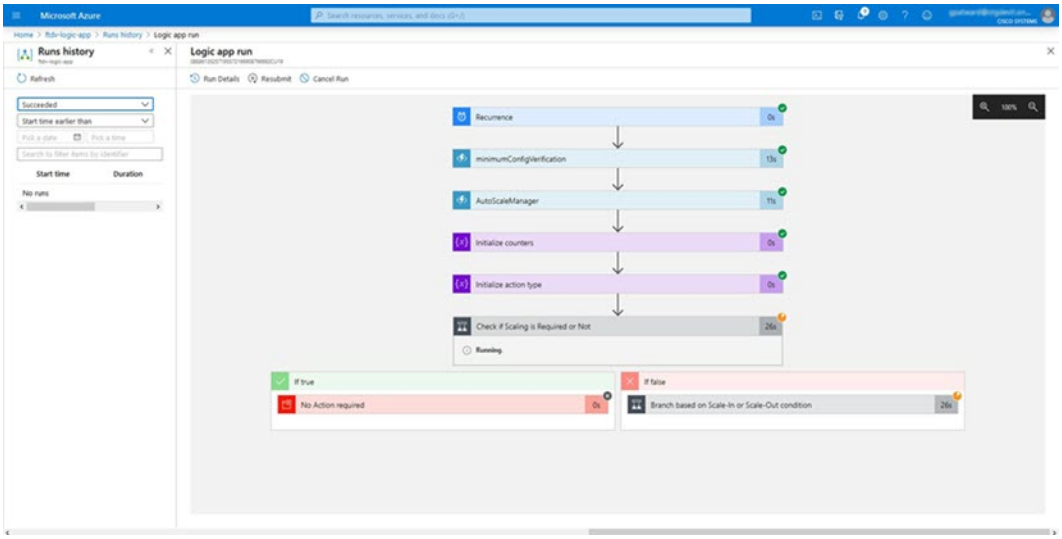

Step 5 Once the Logic App starts, all the deployment-related steps are complete.

Step 6 Verify in the VMSS that Firewall Threat Defense Virtual instances are being created.

Figure 284: Threat Defense Virtual Instances Running

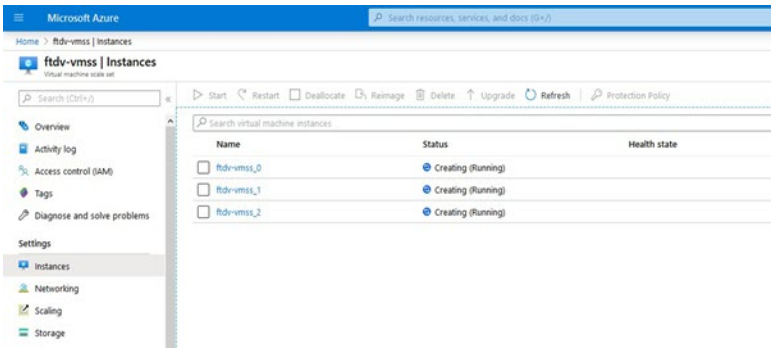

In this example, three Firewall Threat Defense Virtual instances are launched because 'minFtdCount' was set to '3' and 'initDeploymentMode' was set to 'BULK' in the ARM template deployment.

Deploy the Cluster in GCP

To deploy a cluster in GCP, you can either manually deploy or use an instance template to deploy an instance group. You can use the cluster with native GCP load-balancers, or non-native load balancers such as the Cisco Cloud Services Router.

Note Outbound traffic requires interface NAT and is limited to 64K connections.

Sample Topology

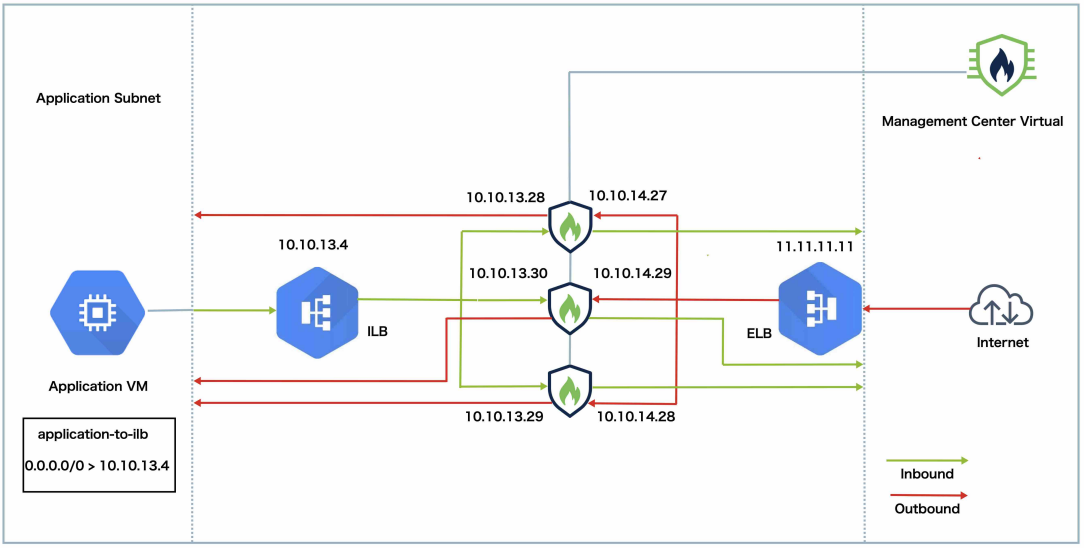

This topology depicts both inbound and outbound traffic flow. The Threat Defense Virtual cluster is sandwiched between the internal and external load balancers. A Management Center Virtual instance is used to manage the cluster.

Inbound traffic from the internet goes to the external load balancer which then transmits the traffic to the Threat Defense Virtual cluster. After the traffic has been inspected by a Threat Defense Virtual instance in the cluster, it is forwarded to the application VM.

Outbound traffic from the application VM is transmitted to the internal load balancer. Traffic is then forwarded to the Threat Defense Virtual cluster and then sent out to the internet.

End-to-End Process for Deploying Threat Defense Virtual Cluster in GCP

Template-based Deployment

The following flowchart illustrates the workflow for template-based deployment of the Threat Defense Virtual cluster on GCP.

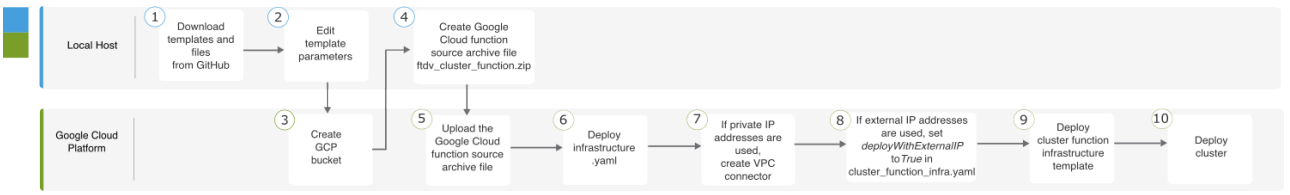

	Workspace	Steps
1	Local Host	Download templates and files from GitHub.
2	Local Host	Edit template parameters.
3	Google Cloud Platform	Create GCP bucket.
4	Local Host	Create Google Cloud function source archive file *ftdv_cluster_function.zip*.
5	Google Cloud Platform	Upload the Google function source archive file.
6	Google Cloud Platform	Deploy infrastructure.yaml.
7	Google Cloud Platform	If private IP addresses are used, create VPC connector.
8	Google Cloud Platform	If external IP addresses are used, set *deployWithExternalIP* to *True* in *cluster_function_infra.yaml*.
9	Google Cloud Platform	Deploy cluster function infrastructure template.
10	Google Cloud Platform	Deploy cluster.

Manual Deployment

The following flowchart illustrates the workflow for manual deployment of the Threat Defense Virtual cluster on GCP.

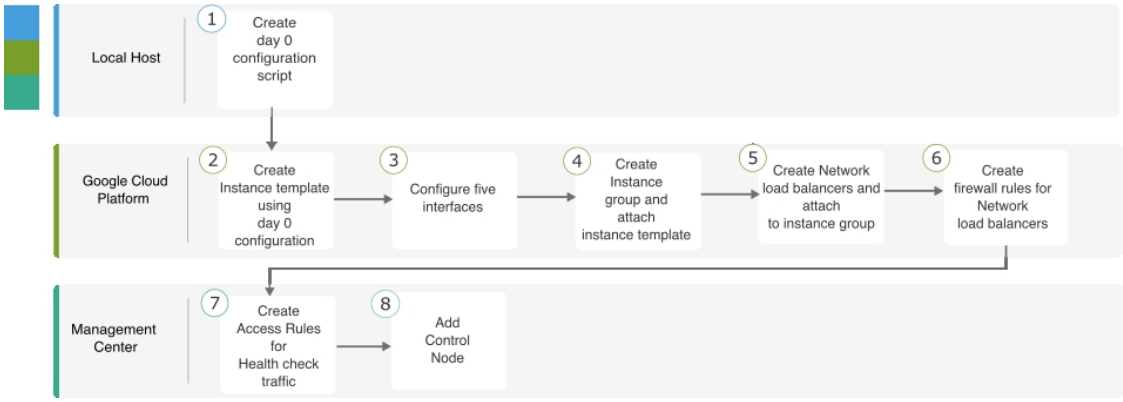

	Workspace	Steps
1	Local Host	Create day 0 configuration script.
2	Google Cloud Platform	Create instance template using day 0 configuration.

	Workspace	Steps
3	Google Cloud Platform	Configure the interfaces.
4	Google Cloud Platform	Create instance group and attach instance template.
5	Google Cloud Platform	Create NLB and attach to instance group.
6	Google Cloud Platform	Create firewall rules for NLB.
7	Management Center	Create access rules for health check traffic.
8	Management Center	Add control node.

Templates

The templates given below are available in GitHub. The parameter values are self-explanatory with the parameter names, and values, given in the template.

- Cluster deployment template for East-West traffic - *deploy_ngfw_cluster_yaml*
- Cluster deployment template for North-South traffic - deploy_ngfw_cluster.yaml

Deploy the Instance Group in GCP Using an Instance Template

Deploy the instance group in GCP using an instance template.

Before you begin

- Use Google Cloud Shell for deployment. Alternatively, you can use Google SDK on any macOS/Linux/Windows machine.
- To allow the cluster to auto-register with the Management Center, you need to create a user with administrative privileges on the Management Center that can use the REST API. See the Cisco Secure Firewall Management Center Administration Guide.
- Add an access policy in the Management Center that matches the name of the policy that you specified in *cluster_function_infra.yaml*.

Procedure

Step 1 Download the templates from GitHub to your local folder.

Step 2 Edit **infrastructure.yaml** , **cluster_function_infra.yaml** and **deploy_ngfw_cluster.yaml** with the required *resourceNamePrefix* parameter (for example, ngfwvcls) and other required user inputs.

From Secure Firewall version 7.4.1, you can deploy the cluster without the diagnostic interface. To deploy the cluster with only the Outside, Inside, Management, and CCL interfaces, set the *withDiagnostic* variable to **False** in both the **infrastructure.yaml** and the **deploy_ngfw_cluster.yaml** files.

Note that there is a **deploy_ngfw_cluster.yaml** file in both the **east-west** and **north-south** folders in GitHub. Download the appropriate template as per your traffic flow requirement.

Step 3 Create a bucket using Google Cloud Shell to upload the Google cloud function source archive file *ftdv_cluster_function.zip*.

gsutil mb --pap enforced gs://*resourceNamePrefix***-ftdv-cluster-bucket/**

Ensure that the *resourceNamePrefix* variable here matches the *resourceNamePrefix* variable that you specified in **cluster_function_infra.yaml**.

Step 4 Create an archive file for the cluster infrastructure.

Example:

```
zip -j ftdv_cluster_function.zip ./cluster-function/*
```

Step 5 Upload the Google source archive that you created earlier.

gsutil cp ftdv_cluster_function.zip gs://*resourceNamePrefix***-ftdv-cluster-bucket/**

Step 6 Deploy infrastructure for the cluster.

gcloud deployment-manager deployments create *cluster_name* **--config infrastructure.yaml**

Step 7 If you are using private IP addresses, perform the steps given below:

a) Launch and set up the Management Center Virtual with a Threat Defense Virtual management VPC.

b) Create a VPC connector to connect the Google Cloud functions with the Threat Defense Virtual management VPC.

gcloud compute networks vpc-access connectors create *vpc-connector-name* **--region us-central1 --subnet** *resourceNamePrefix***-ftdv-mgmt-subnet28**

Step 8 If the Management Center is remote from the Threat Defense Virtual, and the Threat Defense Virtual needs an external IP address, ensure that you set **deployWithExternalIP** to **True** in **cluster_function_infra.yaml**.

Step 9 Deploy the cluster function infrastructure.

gcloud deployment-manager deployments create *cluster_name* **--config cluster_function_infra.yaml**

Step 10 Deploy the cluster.

a. For North-South topology deployment:

gcloud deployment-manager deployments create *cluster_name* **--config north-south/deploy_ngfw_cluster.yaml**

b. For East-West topology deployment:

gcloud deployment-manager deployments create *cluster_name* **--config east-west/deploy_ngfw_cluster.yaml**

Deploy the Cluster in GCP Manually

To deploy the cluster manually, prepare the day0 configuration, deploy each node, and then add the control node to the Firewall Management Center.

Create the Day0 Configuration for GCP

You can use either a fixed configuration or a customized configuration.

Create the Day0 Configuration With a Fixed Configuration for GCP

The fixed configuration will auto-generate the cluster bootstrap configuration.

```
{
    "AdminPassword": "password",
    "Hostname": "hostname",
    "FirewallMode": "Routed",
    "ManageLocally": "No",
                "Diagnostic": "OFF",       //Optional user input from version 7.4.1 - use
to deploy cluster without Diagnostic interface
    "Cluster": {
        "CclSubnetRange": "ip_address_start ip_address_end",
        "ClusterGroupName": "cluster_name"
    }
}
```

For example:

```
{
    "AdminPassword": "DeanW1nche$ter",
    "Hostname": "ciscoftdv",
    "FirewallMode": "Routed",
    "ManageLocally": "No",
    "Cluster": {
        "CclSubnetRange": "10.10.55.2 10.10.55.253",     //mandatory user input
        "ClusterGroupName": "ftdv-cluster"               //mandatory user input
    }
}
```

Note If you are copying and pasting the configuration given above, ensure that you remove **//mandatory user input** from the configuration.

For the **CclSubnetRange** variable, note that you cannot use the first two IP addresses and the last two IP addresses in the subnet. See Reserved IP addresses in IPv4 subnets for more information. Ensure that you have at least 16 available IP addresses for clustering. Some examples of start and end IP addresses are given below.

Table 41: Examples of Start and End IP addresses

CIDR	Start IP Address	End IP Address
10.1.1.0/27	10.1.1.2	10.1.1.29
10.1.1.32/27	10.1.1.34	10.1.1.61

CIDR	Start IP Address	End IP Address
10.1.1.64/27	10.1.1.66	10.1.1.93
10.1.1.96/27	10.1.1.98	10.1.1.125
10.1.1.128/27	10.1.1.130	10.1.1.157
10.1.1.160/27	10.1.1.162	10.1.1.189
10.1.1.192/27	10.1.1.194	10.1.1.221
10.1.1.224/27	10.1.1.226	10.1.1.253
10.1.1.0/24	10.1.1.2	10.1.1.253

Create the Day0 Configuration With a Customized Configuration for GCP

You can enter the entire cluster bootstrap configuration using commands.

```
{
    "AdminPassword": "password",
    "Hostname": "hostname",
    "FirewallMode": "Routed",
    "ManageLocally": "No",
    "run_config": [comma_separated_threat_defense_configuration]
    }
```

The following example creates a configuration with Management, Inside, and Outside interfaces, and a VXLAN interface for the cluster control link. Note the values in bold that need to be unique per node.

```
{
 "AdminPassword": "W1nch3sterBr0s",
 "Hostname": "ftdv1",
 "FirewallMode": "Routed",
 "ManageLocally": "No",
 "run_config": [
  "cluster interface-mode individual force",
  "interface Management0/0",
  "management-only",
  "nameif management",
  "ip address dhcp",
  "interface GigabitEthernet0/0",
  "no shutdown",
  "nameif outside",
  "ip address dhcp",
  "interface GigabitEthernet0/1",
  "no shutdown",
  "nameif inside",
  "ip address dhcp",
  "interface GigabitEthernet0/2",
  "nve-only cluster",
  "nameif ccl_link",
  "ip address dhcp",
  "no shutdown",
  "interface vni1",
  "description Clustering Interface",
  "segment-id 1",
  "vtep-nve 1",
  "object network ccl#link",
  "range 10.1.90.2 10.1.90.17",
  "object-group network cluster#group",
```

```
  "network-object object ccl#link",
  "nve 1",
  "encapsulation vxlan",
  "source-interface ccl_link",
  "peer-group cluster#group",
  "cluster group ftdv-cluster",
  "local-unit 1",
  "cluster-interface vni1 ip 10.1.1.1 255.255.255.0",
  "priority 1",
  "enable",
  "mtu outside 1400",
  "mtu inside 1400"
 ]
}
```

Note For the cluster control link network object, specify only as many addresses as you need (up to 16). A larger range can affect performance.

Deploy Cluster Nodes Manually

Deploy the cluster nodes so they form a cluster. For clustering on GCP, you cannot use the 4 vCPU machine type. The 4 vCPU machine type only supports four interfaces, and five are needed. Use a machine type that supports five interfaces, such as c2-standard-8.

Procedure

Step 1 Create an instance template using the day 0 configuration (in the **Metadata > Startup Script** section) with 5 interfaces: outside, inside, management, diagnostic, and cluster control link.

See Cisco Secure Firewall Threat Defense Virtual Getting Started Guide.

Step 2 Create an instance group, and attach the instance template.

Step 3 Create GCP network load balancers (internal and external), and attach the instance group.

Step 4 For GCP network load balancers, allow health checks in your security policy on the Management Center. See Allow Health Checks for GCP Network Load Balancers, on page 660.

Step 5 Add the control node to the Management Center. See Add the Cluster to the Management Center (Manual Deployment), on page 662.

Allow Health Checks for GCP Network Load Balancers

Google Cloud provides health checks to determine if backends respond to traffic.

See https://cloud.google.com/load-balancing/docs/health-checks to create firewall rules for network load balancers. Then in the Firewall Management Center, create access rules to allow the health check traffic. See https://cloud.google.com/load-balancing/docs/health-check-concepts for the required network ranges. See Access Control Rules, on page 1773.

You also need to configure dynamic manual NAT rules to redirect the health check traffic to the Google metadata server at 169.254.169.254. See Configure Dynamic Manual NAT, on page 1041.

You can set up a route for GCP health checks across all interfaces that are used to configure their health probes. You can achieve this by creating a route with a higher metric on interfaces where a route for GCP health checks is not already available.

North-South NAT Rules Sample Configuration

```
nat (inside,outside) source dynamic GCP-HC ILB-SOUTH destination static ILB-SOUTH METADATA
nat (outside,outside) source dynamic GCP-HC ELB-NORTH destination static ELB-NORTH METADATA

nat (outside,inside) source static any interface destination static ELB-NORTH Ubuntu-App-VM
nat (inside,outside) source dynamic any interface destination static obj-any obj-any

object network Metadata
 host 169.254.169.254

object network ILB-SOUTH
host <ILB_IP>
object network ELB-NORTH
host <ELB_IP>

object-group network GCP-HC
 network-object 35.191.0.0 255.255.0.0
 network-object 130.211.0.0 255.255.252.0
 network-object 209.85.204.0 255.255.252.0
 network-object 209.85.152.0 255.255.252.0
```

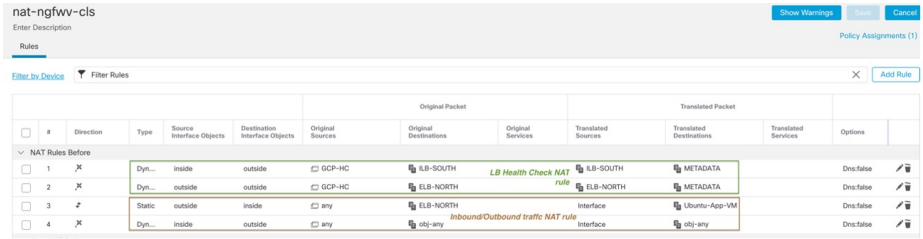

East-West NAT Rules Sample Configuration

```
nat (inside,outside) source dynamic GCP-HC ILB-East destination static ILB-East Metadata
nat (outside,outside) source dynamic GCP-HC ILB-West destination static ILB-West Metadata

object network Metadata
 host 169.254.169.254

object network ILB-East
host <ILB_East_IP>
object network ILB-West
host <ILB_West_IP>

object-group network GCP-HC
 network-object 35.191.0.0 255.255.0.0
 network-object 130.211.0.0 255.255.252.0
 network-object 209.85.204.0 255.255.252.0
 network-object 209.85.152.0 255.255.252.0
```

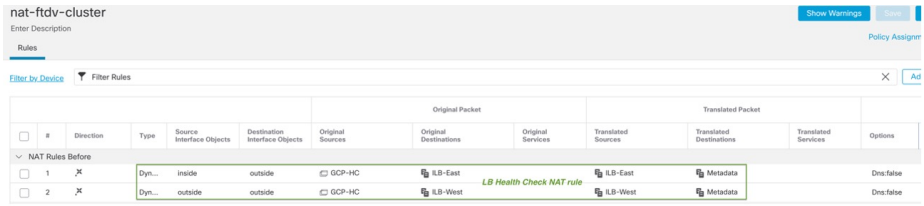

North-South and East-West Traffic Routing Configuration Sample

```
route outside 0.0.0.0 0.0.0.0 <Outside_Gateway> 1
route inside 35.191.0.0 255.255.0.0 <Inside_Gateway> 1
route inside 130.211.0.0 255.255.252.0 <Inside_Gateway> 1
route inside 209.85.152.0 255.255.252.0 <Inside_Gateway> 1
route inside 209.85.204.0 255.255.252.0 <Inside_Gateway> 1
```

If a default route is not available, then policy-based routing can be used to route the traffic for health checks.

Add the Cluster to the Management Center (Manual Deployment)

Use this procedure to add the cluster to the Firewall Management Center if you manually deployed the cluster. If you used a template, the cluster will auto-register on the Firewall Management Center.

Add one of the cluster units as a new device to the Firewall Management Center; the Firewall Management Center auto-detects all other cluster members.

Before you begin

- All cluster units must be in a successfully-formed cluster prior to adding the cluster to the Firewall Management Center. You should also check which unit is the control unit. Use the Firewall Threat Defense **show cluster info** command.

Procedure

Step 1 In the Firewall Management Center, choose **Devices** > **Device Management**, and then choose **Add** > **Add Device** to add the control unit using the unit's management IP address.

Figure 285: Add Device

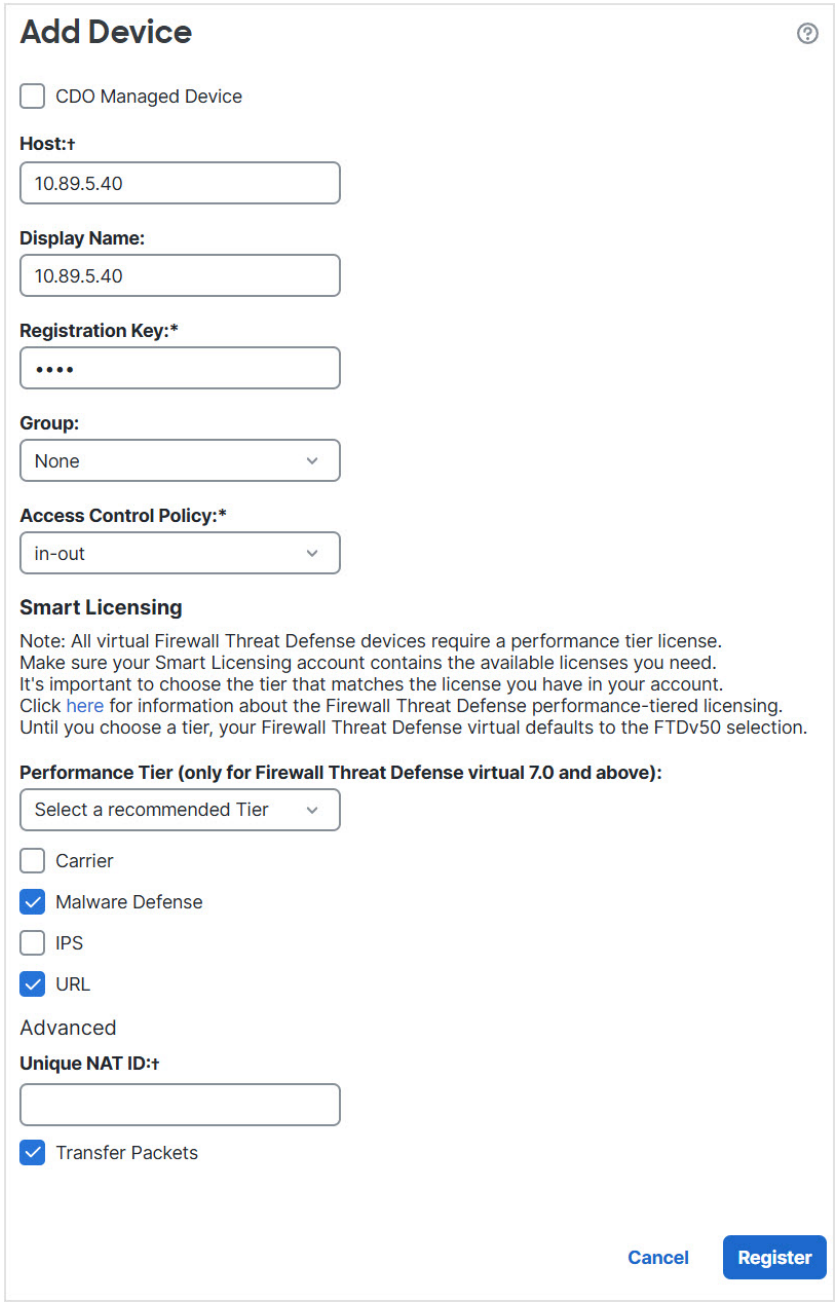

a) In the **Host** field, enter the IP address or hostname of the control unit.

We recommend adding the control unit for the best performance, but you can add any unit of the cluster.

If you used a NAT ID during device setup, you may not need to enter this field. For more information, see NAT Environments, on page 7.

b) In the **Display Name** field, enter a name for the control unit as you want it to display in the Firewall Management Center.

This display name is not for the cluster; it is only for the control unit you are adding. You can later change the name of other cluster members and the cluster display name.

c) In the **Registration Key** field, enter the same registration key that you used during device setup. The registration key is a one-time-use shared secret.
d) (Optional) Add the device to a device **Group**.
e) Choose an initial **Access Control Policy** to deploy to the device upon registration, or create a new policy.

If you create a new policy, you create a basic policy only. You can later customize the policy as needed.

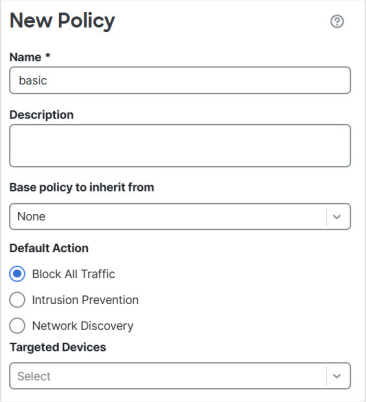

f) Choose licenses to apply to the device.
g) If you used a NAT ID during device setup, expand the **Advanced** section and enter the same NAT ID in the **Unique NAT ID** field.
h) Check the **Transfer Packets** check box to allow the device to transfer packets to the Firewall Management Center.

This option is enabled by default. When events like IPS or Snort are triggered with this option enabled, the device sends event metadata information and packet data to the Firewall Management Center for inspection. If you disable it, only event information will be sent to the Firewall Management Center but packet data is not sent.

i) Click **Register**.

The Firewall Management Center identifies and registers the control unit, and then registers all data units. If the control unit does not successfully register, then the cluster is not added. A registration failure can occur if the cluster was not up, or because of other connectivity issues. In this case, we recommend that you try re-adding the cluster unit.

The cluster name shows on the **Devices** > **Device Management** page; expand the cluster to see the cluster units.

Figure 286: Cluster Management

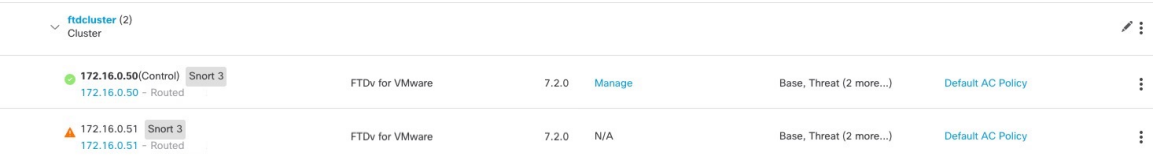

A unit that is currently registering shows the loading icon.

Figure 287: Node Registration

Note
GCP prioritizes nodes with public IP address during cluster node discovery. To ensure the Firewall Threat Defense Virtual cluster registers with the management center virtual using the private IP address, you must first disable the public IP address on the Firewall Threat Defense Virtual cluster node. This allows GCP node discovery to proceed using the private IP address for registration node with the management center virtual.

You can monitor cluster unit registration by clicking the **Notifications** icon and choosing **Tasks**. The Firewall Management Center updates the Cluster Registration task as each unit registers. If any units fail to register, see Reconcile Cluster Nodes, on page 674.

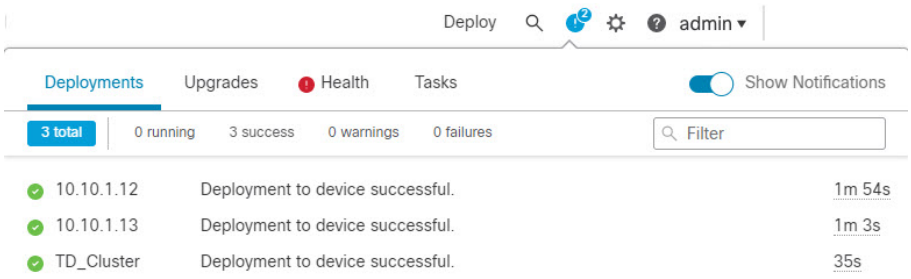

Step 2 Configure device-specific settings by clicking the **Edit** (✎) for the cluster.

Most configuration can be applied to the cluster as a whole, and not nodes in the cluster. For example, you can change the display name per node, but you can only configure interfaces for the whole cluster.

Step 3 On the **Devices** > **Device Management** > **Cluster** screen, you see **General**, **License**, **System**, and **Health** settings.

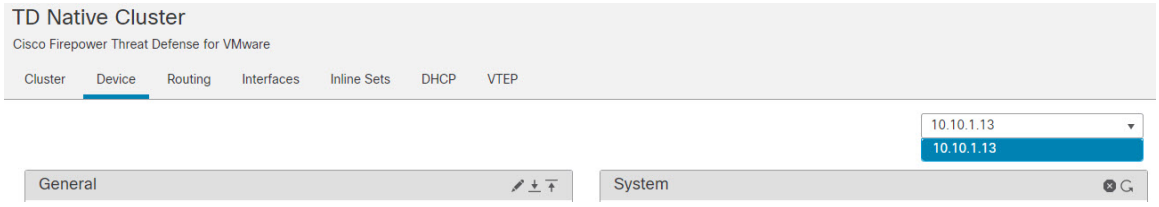

See the following cluster-specific items:

- **General** > **Name**—Change the cluster display name by clicking the **Edit** (✎).

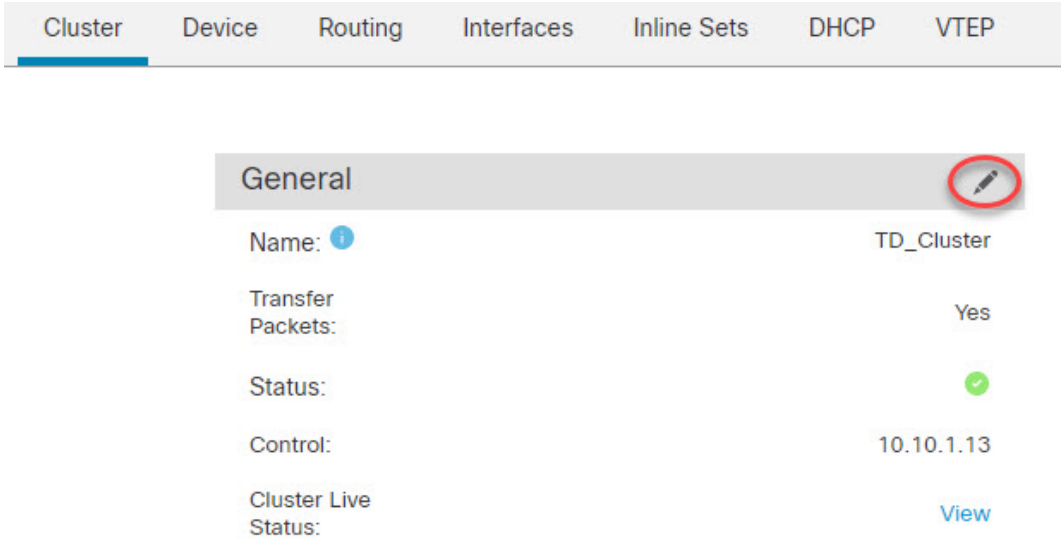

Then set the **Name** field.

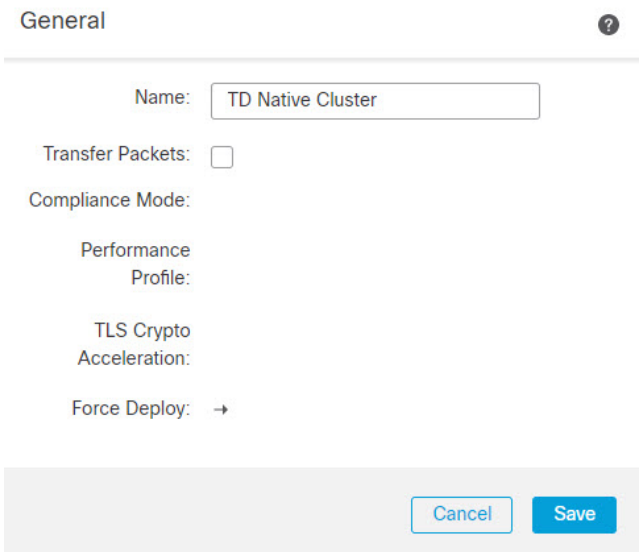

- **General > Cluster Live Status**—Click the **View** link to open the **Cluster Status** dialog box.

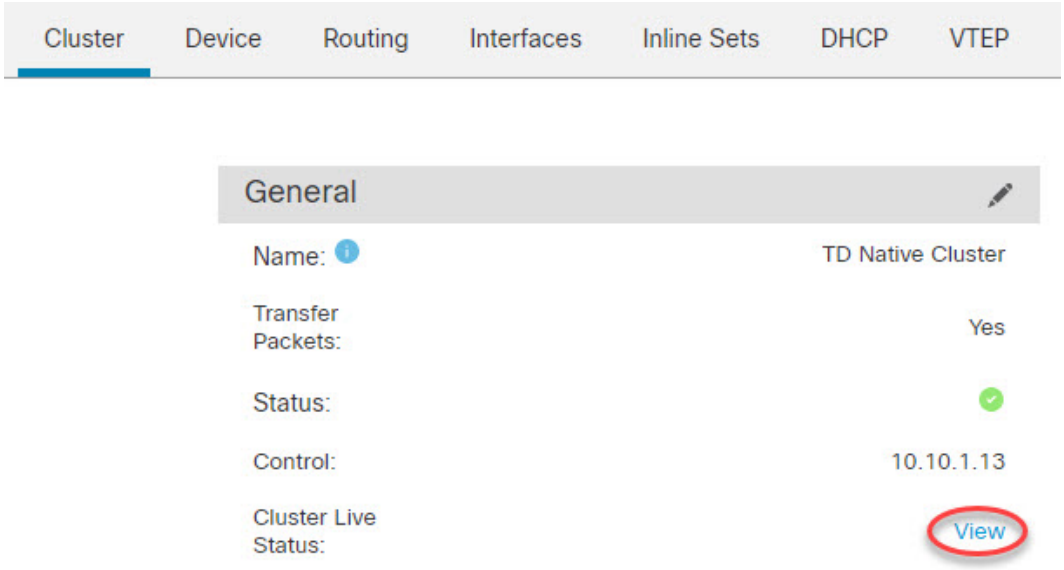

The **Cluster Status** dialog box also lets you retry data unit registration by clicking **Reconcile**.You can also ping the cluster control link from a node. See Perform a Ping on the Cluster Control Link, on page 682.

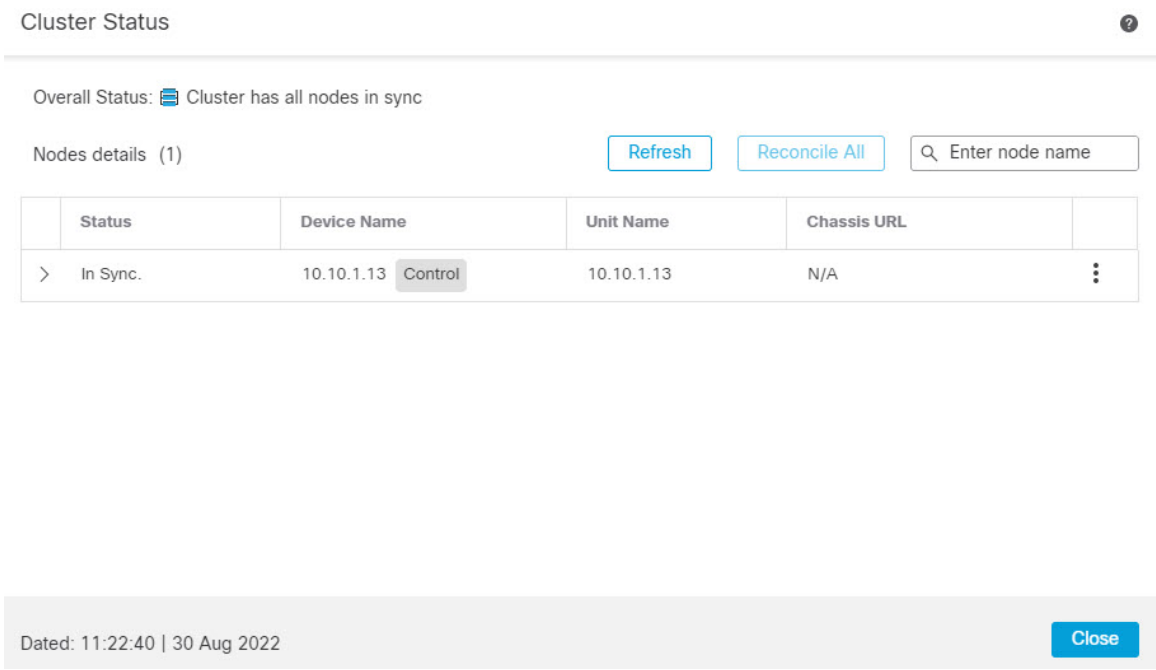

- **General > Troubleshoot**—You can generate and download troubleshooting logs, and you can view cluster CLIs. See Troubleshooting the Cluster, on page 682.

Figure 288: Troubleshoot

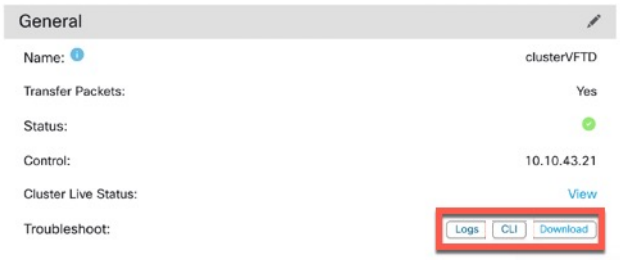

- **License**—Click **Edit** (✎) to set license entitlements.

Step 4 On the **Devices** > **Device Management** > **Devices**, you can choose each member in the cluster from the top right drop-down menu and configure the following settings.

- **General > Name**—Change the cluster member display name by clicking the **Edit** (✎).

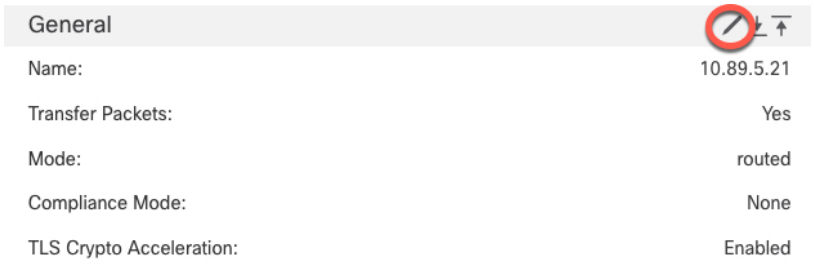

Then set the **Name** field.

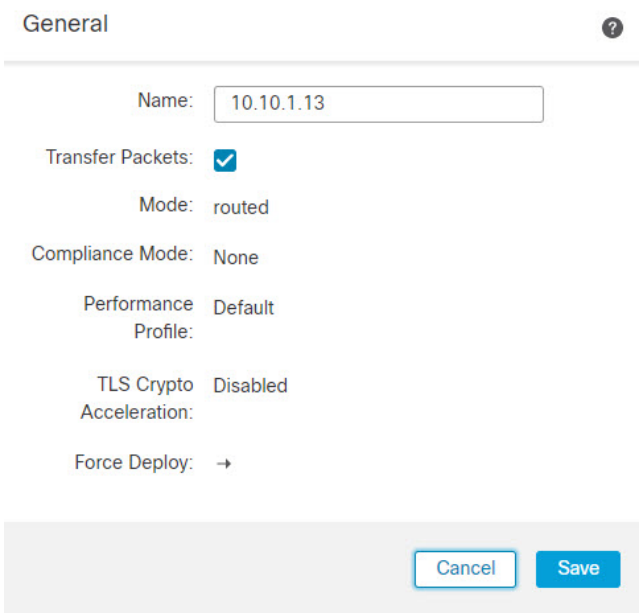

- **Management > Host**—If you change the management IP address in the device configuration, you must match the new address in the Firewall Management Center so that it can reach the device on the network; edit the **Host** address in the **Management** area.

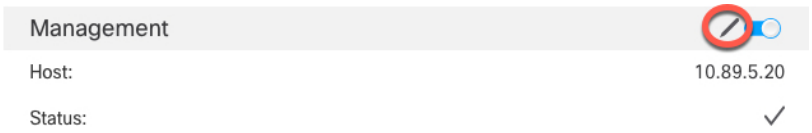

Configure Cluster Health Monitor Settings

The **Cluster Health Monitor Settings** section of the **Cluster** page displays the settings described in the table below.

Figure 289: Cluster Health Monitor Settings

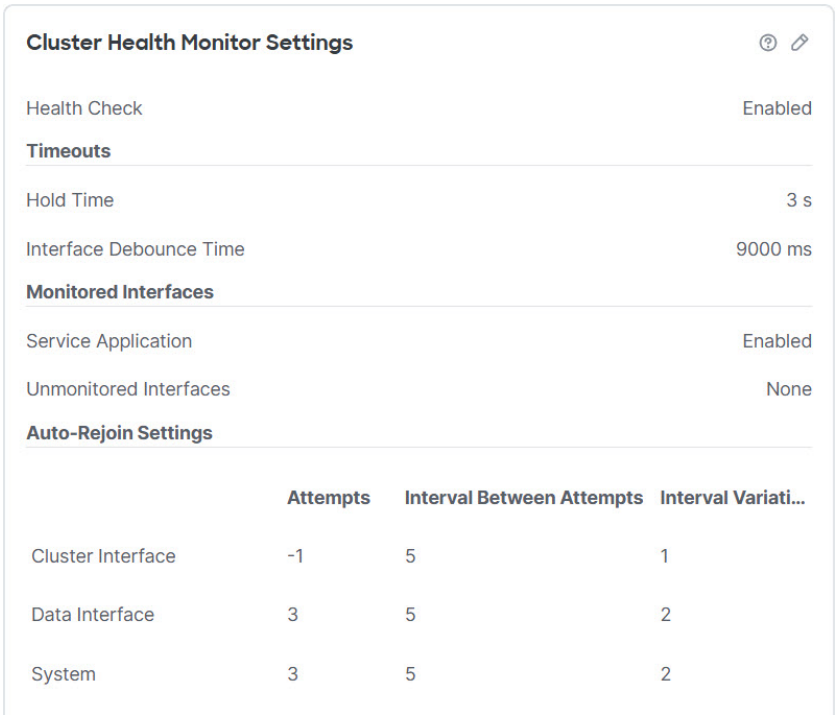

Table 42: Cluster Health Monitor Settings Section Table Fields

Field	Description
Timeouts	

Field	Description
Hold Time	Between .3 and 45 seconds; The default is 3 seconds. To determine node system health, the cluster nodes send heartbeat messages on the cluster control link to other nodes. If a node does not receive any heartbeat messages from a peer node within the hold time period, the peer node is considered unresponsive or dead.
Interface Debounce Time	Between 300 and 9000 ms. The default is 500 ms. The interface debounce time is the amount of time before the node considers an interface to be failed, and the node is removed from the cluster.
Monitored Interfaces	The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The amount of time before the node removes a member from the cluster depends on the type of interface and whether the node is an established node or is joining the cluster.
Service Application	Shows whether the Snort and disk-full processes are monitored.
Unmonitored Interfaces	Shows unmonitored interfaces.
Auto-Rejoin Settings	
Cluster Interface	Shows the auto-rejoin settings after a cluster control link failure.
Attempts	Between -1 and 65535. The default is -1 (unlimited). Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 1x the interval duration. Defines if the interval duration increases at each attempt.
Data Interfaces	Shows the auto-rejoin settings after a data interface failure.
Attempts	Between -1 and 65535. The default is 3. Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 2x the interval duration. Defines if the interval duration increases at each attempt.
System	Shows the auto-rejoin settings after internal errors. Internal failures include: application sync timeout; inconsistent application statuses; and so on.
Attempts	Between -1 and 65535. The default is 3. Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.

Field	Description
Interval Variation	Between 1 and 3. The default is 2x the interval duration. Defines if the interval duration increases at each attempt.

Note If you disable the system health check, fields that do not apply when the system health check is disabled will not show.

You can change these settings from this section.

You can monitor any port-channel ID, single physical interface ID, as well as the Snort and disk-full processes. Health monitoring is not performed on VLAN subinterfaces or virtual interfaces such as VNIs or BVIs. You cannot configure monitoring for the cluster control link; it is always monitored.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the cluster you want to modify, click **Edit** (✎).

Step 3 Click **Cluster**.

Step 4 In the **Cluster Health Monitor Settings** section, click **Edit** (✎).

Step 5 Disable the system health check by clicking the **Health Check** slider .

Figure 290: Disable the System Health Check

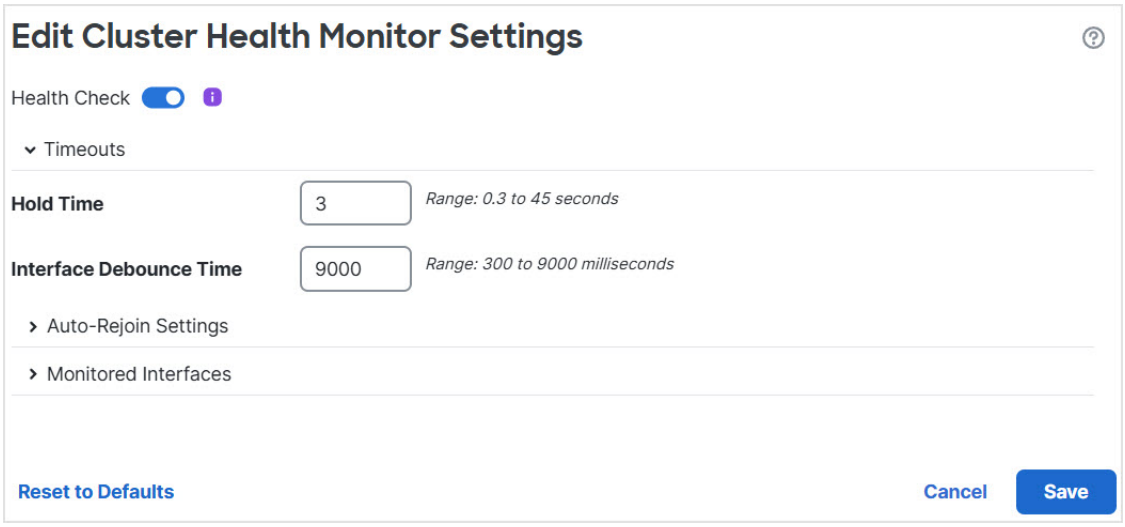

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 6 Configure the hold time and interface debounce time.

- **Hold Time**—Set the hold time to determine the amount of time between node heartbeat status messages, between .3 and 45 seconds; The default is 3 seconds.
- **Interface Debounce Time**—Set the debounce time between 300 and 9000 ms. The default is 500 ms. Lower values allow for faster detection of interface failures. Note that configuring a lower debounce time increases the chances of false-positives. When an interface status update occurs, the node waits the number of milliseconds specified before marking the interface as failed, and the node is removed from the cluster. In the case of an EtherChannel that transitions from a down state to an up state (for example, the switch reloaded, or the switch enabled an EtherChannel), a longer debounce time can prevent the interface from appearing to be failed on a cluster node just because another cluster node was faster at bundling the ports.

Step 7 Customize the auto-rejoin cluster settings after a health check failure.

Figure 291: Configure Auto-Rejoin Settings

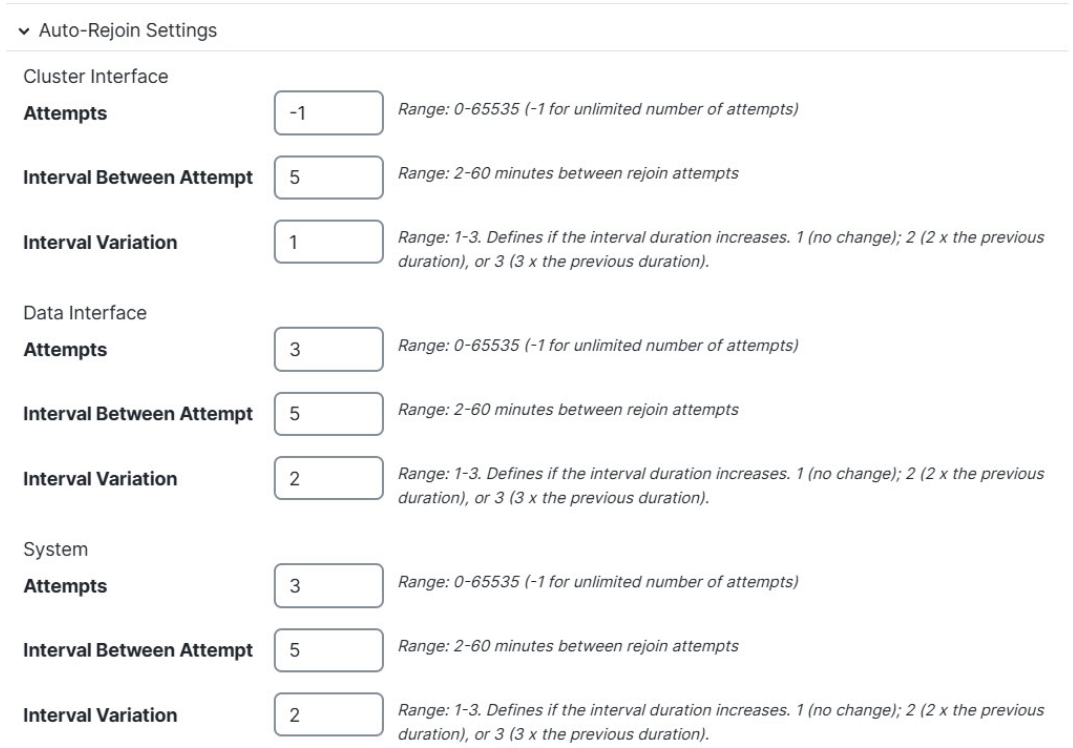

Set the following values for the **Cluster Interface**, **Data Interface**, and **System** (internal failures include: application sync timeout; inconsistent application statuses; and so on):

- **Attempts**—Sets the number of rejoin attempts, between -1 and 65535. **0** disables auto-rejoining. The default for the **Cluster Interface** is -1 (unlimited). The default for the **Data Interface** and **System** is 3.
- **Interval Between Attempts**—Defines the interval duration in minutes between rejoin attempts, between 2 and 60. The default value is 5 minutes. The maximum total time that the node attempts to rejoin the cluster is limited to 14400 minutes (10 days) from the time of last failure.

- **Interval Variation**—Defines if the interval duration increases. Set the value between 1 and 3: **1** (no change); **2** (2 x the previous duration), or **3** (3 x the previous duration). For example, if you set the interval duration to 5 minutes, and set the variation to 2, then the first attempt is after 5 minutes; the 2nd attempt is 10 minutes (2 x 5); the 3rd attempt 20 minutes (2 x 10), and so on. The default value is **1** for the **Cluster Interface** and **2** for the **Data Interface** and **System**.

Step 8 Configure monitored interfaces by moving interfaces in the **Monitored Interfaces** or **Unmonitored Interfaces** window. You can also check or uncheck **Enable Service Application Monitoring** to enable or disable monitoring of the Snort and disk-full processes.

Figure 292: Configure Monitored Interfaces

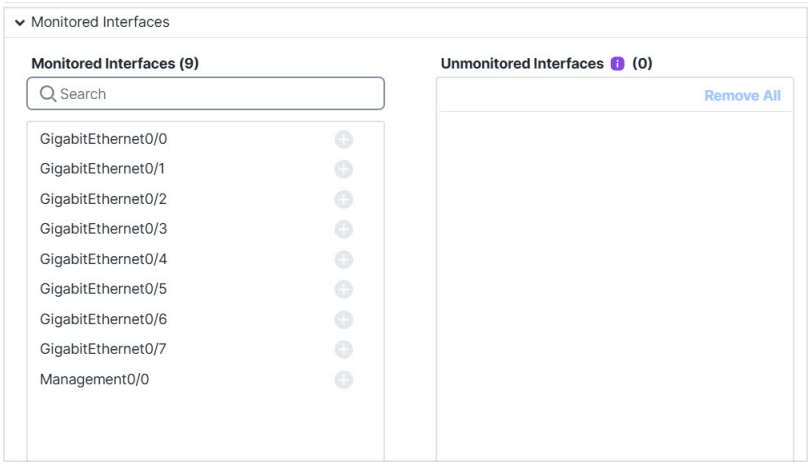

The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The amount of time before the node removes a member from the cluster depends on the type of interface and whether the node is an established node or is joining the cluster. Health check is enabled by default for all interfaces and for the Snort and disk-full processes.

You might want to disable health monitoring of non-essential interfaces.

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 9 Click **Save**.

Step 10 Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Manage Cluster Nodes

Disable Clustering

You may want to deactivate a node in preparation for deleting the node, or temporarily for maintenance. This procedure is meant to temporarily deactivate a node; the node will still appear in the Firewall Management Center device list. When a node becomes inactive, all data interfaces are shut down.

Note Do not power off the node without first disabling clustering.

Procedure

Step 1 For the unit you want to disable, choose **Devices** > **Device Management**, click the **More** (⋮), and choose **Disable Node Clustering**.

Step 2 Confirm that you want to disable clustering on the node.

The node will show **(Disabled)** next to its name in the **Devices** > **Device Management** list.

Step 3 To reenable clustering, see Rejoin the Cluster, on page 674.

Rejoin the Cluster

If a node was removed from the cluster, for example for a failed interface or if you manually disabled clustering, you must manually rejoin the cluster. Make sure the failure is resolved before you try to rejoin the cluster. See Rejoining the Cluster, on page 691 for more information about why a node can be removed from a cluster.

Procedure

Step 1 For the unit you want to reactivate, choose **Devices** > **Device Management**, click the **More** (⋮), and choose **Enable Node Clustering**.

Step 2 Confirm that you want to enable clustering on the node.

Reconcile Cluster Nodes

If a cluster node fails to register, you can reconcile the cluster membership from the device to the Firewall Management Center. For example, a data node might fail to register if the Firewall Management Center is occupied with certain processes, or if there is a network issue.

Procedure

Step 1 Choose **Devices** > **Device Management More** (⋮) for the cluster, and then choose **Cluster Live Status** to open the **Cluster Status** dialog box.

Step 2 Click **Reconcile All**.

Figure 293: Reconcile All

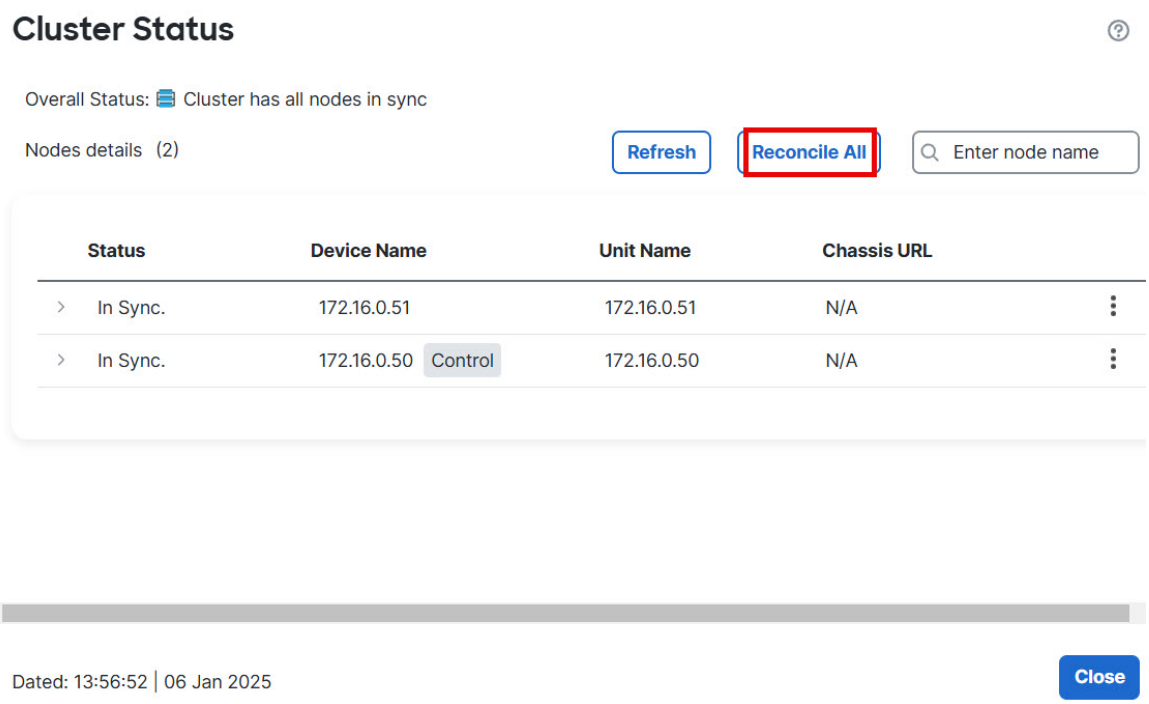

For more information about the cluster status, see Monitoring the Cluster, on page 676.

Unregister the Cluster or Nodes and Register to a New Firewall Management Center

You can unregister the cluster from the Firewall Management Center, which keeps the cluster intact. You might want to unregister the cluster if you want to add the cluster to a new Firewall Management Center.

You can also unregister a node from the Firewall Management Center without breaking the node from the cluster. Although the node is not visible in the Firewall Management Center, it is still part of the cluster, and it will continue to pass traffic and could even become the control node. You cannot unregister the current control node. You might want to unregister the node if it is no longer reachable from the Firewall Management Center, but you still want to keep it as part of the cluster while you troubleshoot management connectivity.

Unregistering a cluster:

- Severs all communication between the Firewall Management Center and the cluster.

- Removes the cluster from the **Device Management** page.
- Returns the cluster to local time management if the cluster's platform settings policy is configured to receive time from the Firewall Management Center using NTP.
- Leaves the configuration intact, so the cluster continues to process traffic.

  Policies, such as NAT and VPN, ACLs, and the interface configurations remain intact.

Registering the cluster again to the same or a different Firewall Management Center causes the configuration to be removed, so the cluster will stop processing traffic at that point; the cluster configuration remains intact so you can add the cluster as a whole. You can choose an access control policy at registration, but you will have to re-apply other policies after registration and then deploy the configuration before it will process traffic again.

Before you begin

This procedure requires CLI access to one of the nodes.

Procedure

Step 1 Choose **Devices** > **Device Management**, click **More** (⋮) for the cluster or node, and choose **Unregister**.

Step 2 You are prompted to unregister the cluster or node; click **Yes**.

Step 3 You can register the cluster to a new (or the same) Firewall Management Center by adding one of the cluster members as a new device.

You only need to add one of the cluster nodes as a device, and the rest of the cluster nodes will be discovered.

a) Connect to one cluster node's CLI, and identify the new Firewall Management Center using the **configure manager add** command.
b) Choose **Devices** > **Device Management**, and then click **Add Device**.

Step 4 To re-add a deleted node, see Reconcile Cluster Nodes, on page 674.

Monitoring the Cluster

You can monitor the cluster in the Firewall Management Center and at the Firewall Threat Defense CLI.

- **Cluster Status** dialog box, which is available from the **Devices** > **Device Management More** (⋮) icon or from the **Devices** > **Device Management** > **Cluster** page **>** **General** area **>** **Cluster Live Status** link.

Figure 294: Cluster Status

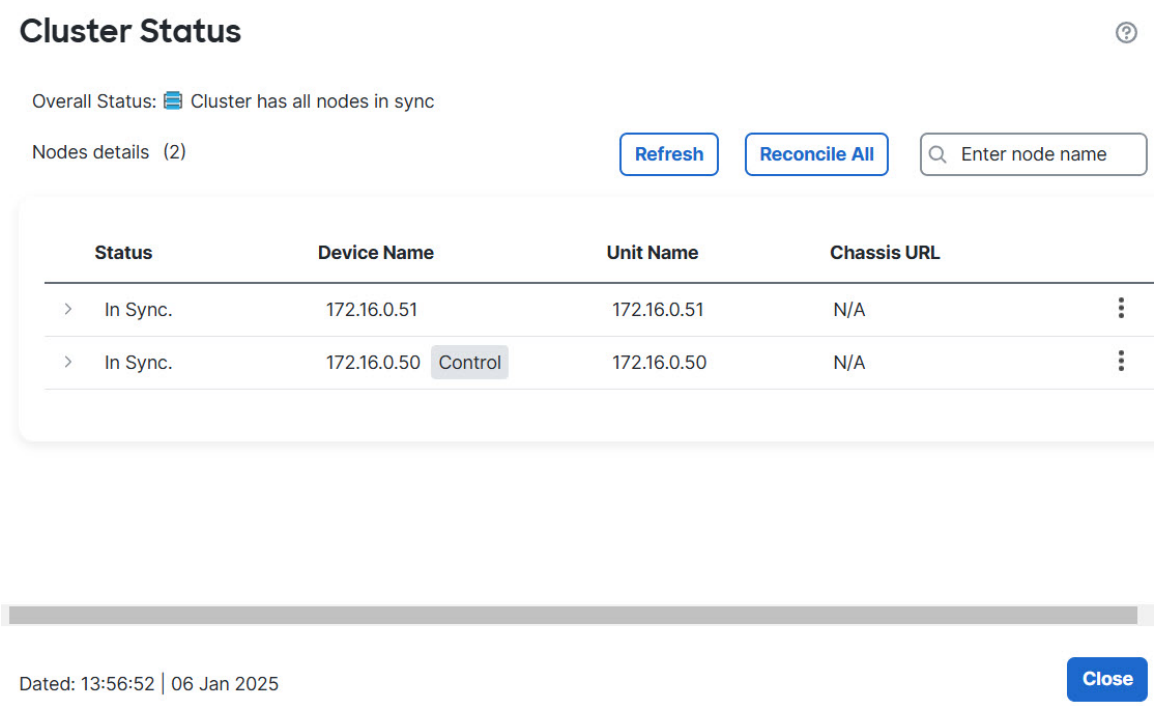

The Control node has a graphic indicator identifying its role.

Cluster member **Status** includes the following states:

- In Sync.—The node is registered with the Firewall Management Center.
- Pending Registration—The node is part of the cluster, but has not yet registered with the Firewall Management Center. If a node fails to register, you can retry registration by clicking **Reconcile All**.
- Clustering is disabled—The node is registered with the Firewall Management Center, but is an inactive member of the cluster. The clustering configuration remains intact if you intend to later re-enable it, or you can delete the node from the cluster.
- Joining cluster...—The node is joining the cluster on the chassis, but has not completed joining. After it joins, it will register with the Firewall Management Center.

For each node, you can view the **Summary** or the **History**.

Figure 295: Node Summary

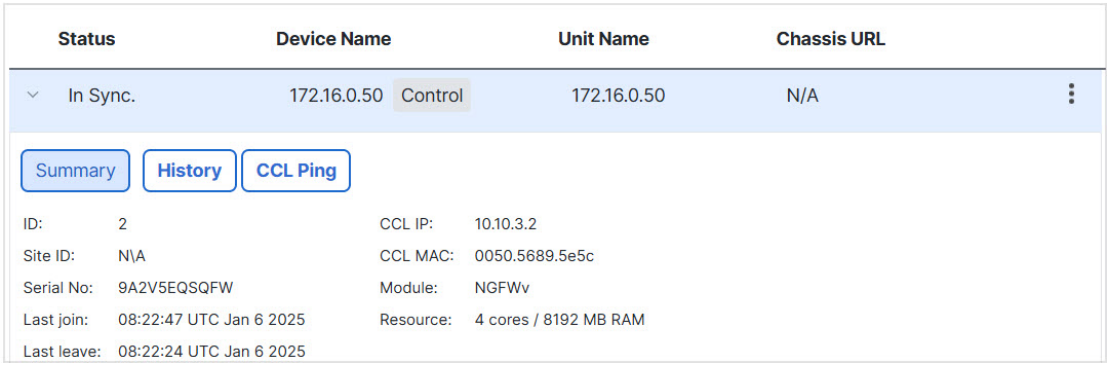

Figure 296: Node History

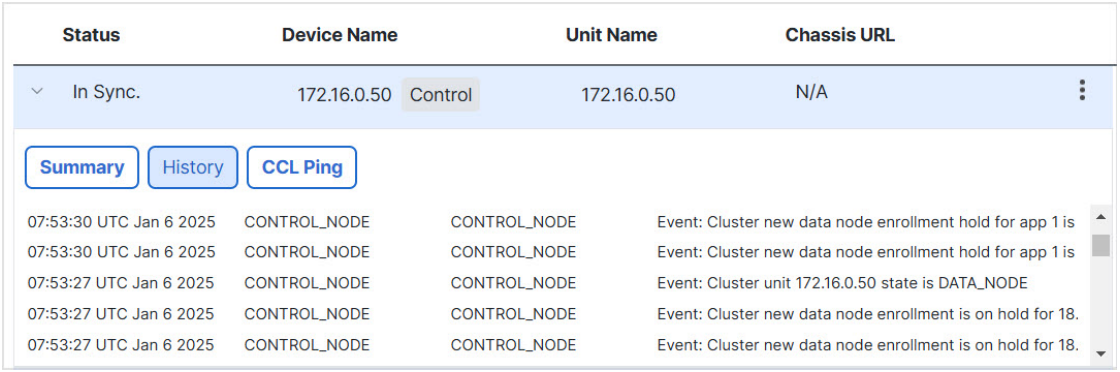

- **System** (⚙) > **Tasks** page.

  The **Tasks** page shows updates of the Cluster Registration task as each node registers.

- **Devices** > **Device Management***cluster_name*.

  When you expand the cluster on the devices listing page, you can see all member nodes, including the control node shown with its role next to the IP address. For nodes that are still registering, you can see the loading icon.

- **show cluster** {**access-list** [*acl_name*] | **conn** [**count**] | **cpu** [**usage**] | **history** | **interface-mode** | **memory** | **resource usage** | **service-policy** | **traffic** | **xlate count**}

  To view aggregated data for the entire cluster or other information, use the **show cluster** command.

- **show cluster info** [**auto-join** | **clients** | **conn-distribution** | **flow-mobility counters** | **goid** [*options*] | **health** | **incompatible-config** | **loadbalance** | **old-members** | **packet-distribution** | **trace** [*options*] | **transport** { **asp** | **cp**}]

  To view cluster information, use the **show cluster info** command.

Cluster Health Monitor Dashboard

Cluster Health Monitor

When a Firewall Threat Defense is the control node of a cluster, the Firewall Management Center collects various metrics periodically from the device metric data collector. The cluster health monitor is comprised of the following components:

- Overview dashboard—Displays information about the cluster topology, cluster statistics, and metric charts:
  - The topology section displays a cluster's live status, the health of individual threat defense, threat defense node type (control node or data node), and the status of the device. The status of the device could be *Disabled* (when the device leaves the cluster), *Added out of box* (in a public cloud cluster, the additional nodes that do not belong to the Firewall Management Center), or *Normal* (ideal state of the node).
  - The cluster statistics section displays current metrics of the cluster with respect to the CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.

Note The CPU and memory metrics display the individual average of the data plane and snort usage.

  - The metric charts, namely, CPU Usage, Memory Usage, Throughput, and Connections, diagrammatically display the statistics of the cluster over the specified time period.

- Load Distribution dashboard—Displays load distribution across the cluster nodes in two widgets:
  - The Distribution widget displays the average packet and connection distribution over the time range across the cluster nodes. This data depicts how the load is being distributed by the nodes. Using this widget, you can easily identify any abnormalities in the load distribution and rectify it.
  - The Node Statistics widget displays the node level metrics in table format. It displays metric data on CPU usage, memory usage, input rate, output rate, active connections, and NAT translations across the cluster nodes. This table view enables you to correlate data and easily identify any discrepancies.

- Member Performance dashboard—Displays current metrics of the cluster nodes. You can use the selector to filter the nodes and view the details of a specific node. The metric data include CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.
- CCL dashboard—Displays, graphically, the cluster control link data namely, the input, and output rate.
- Troubleshooting and Links — Provides convenient links to frequently used troubleshooting topics and procedures.
- Time range—An adjustable time window to constrain the information that appears in the various cluster metrics dashboards and widgets.
- Custom Dashboard—Displays data on both cluster-wide metrics and node-level metrics. However, node selection only applies for the threat defense metrics and not for the entire cluster to which the node belongs.

Viewing Cluster Health

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

The cluster health monitor provides a detailed view of the health status of a cluster and its nodes. This cluster health monitor provides health status and trends of the cluster in an array of dashboards.

Before you begin

- Ensure you have created a cluster from one or more devices in the Firewall Management Center.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Use the Monitoring navigation pane to access node-specific health monitors.

Step 2 In the device list, click **Expand**(>) and **Collapse** (⋁) to expand and collapse the list of managed cluster devices.

Step 3 To view the cluster health statistics, click on the cluster name. The cluster monitor reports health and performance metrics in several predefined dashboards by default. The metrics dashboards include:

- Overview — Highlights key metrics from the other predefined dashboards, including its nodes, CPU, memory, input and output rates, connection statistics, and NAT translation information.
- Load Distribution — Traffic and packet distribution across the cluster nodes.
- Member Performance — Node-level statistics on CPU usage, memory usage, input throughput, output throughput, active connection, and NAT translation.
- CCL — Interface status and aggregate traffic statistics.

You can navigate through the various metrics dashboards by clicking on the labels. For a comprehensive list of the supported cluster metrics, see Cisco Secure Firewall Threat Defense Health Metrics.

Step 4 You can configure the time range from the drop-down in the upper-right corner. The time range can reflect a period as short as the last hour (the default) or as long as two weeks. Select **Custom** from the drop-down to configure a custom start and end date.

Click the refresh icon to set auto refresh to 5 minutes or to toggle off auto refresh.

Step 5 Click on deployment icon for a deployment overlay on the trend graph, with respect to the selected time range.

The deployment icon indicates the number of deployments during the selected time-range. A vertical band indicates the deployment start and end time. For multiple deployments, multiple bands/lines appear. Click on the icon on top of the dotted line to view the deployment details.

Step 6 (For node-specific health monitor) View the **Health Alerts** for the node in the alert notification at the top of page, directly to the right of the device name.

Hover your pointer over the **Health Alerts** to view the health summary of the node. The popup window shows a truncated summary of the top five health alerts. Click on the popup to open a detailed view of the health alert summary.

Step 7 (For node-specific health monitor) The device monitor reports health and performance metrics in several predefined dashboards by default. The metrics dashboards include:

- Overview — Highlights key metrics from the other predefined dashboards, including CPU, memory, interfaces, connection statistics; plus disk usage and critical process information.
- CPU — CPU utilization, including the CPU usage by process and by physical cores.
- Memory — Device memory utilization, including data plane and Snort memory usage.
- Interfaces — Interface status and aggregate traffic statistics.
- Connections — Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
- Snort — Statistics that are related to the Snort process.
- ASP drops — Statistics related to the dropped packets against various reasons.

You can navigate through the various metrics dashboards by clicking on the labels. See Cisco Secure Firewall Threat Defense Health Metrics for a comprehensive list of the supported device metrics.

Step 8 Click the plus sign **Add New Dashboard**(+) in the upper right corner of the health monitor to create a custom dashboard by building your own variable set from the available metric groups.

For cluster-wide dashboard, choose Cluster metric group, and then choose the metric.

Cluster Metrics

The cluster health monitor tracks statistics that are related to a cluster and its nodes, and aggregate of load distribution, performance, and CCL traffic statistics.

Table 43: Cluster Metrics

Metric	Description	Format
CPU	Average of CPU metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Memory	Average of memory metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Data Throughput	Incoming and outgoing data traffic statistics for a cluster.	bytes
CCL Throughput	Incoming and outgoing CCL traffic statistics for a cluster.	bytes
Connections	Count of active connections in a cluster.	number
NAT Translations	Count of NAT translations for a cluster.	number
Distribution	Connection distribution count in the cluster for every second.	number

Metric	Description	Format
Packets	Packet distribution count in the cluster for every second.	number

Troubleshooting the Cluster

You can use the **CCL Ping** tool to make sure the cluster control link is operating correctly. You can also use the following tools that are available for devices and clusters:

- Troubleshooting files—If a node fails to join the cluster, a troubleshooting file is automatically generated. You can also generate and download troubleshooting files from the **Devices** > **Device Management** > **Cluster** > **General** area. See Generate Troubleshooting Files, on page 129.

  You can also generate files from the **Device Management** page by clicking **More** (⋮) and choosing **Troubleshoot Files**.

- CLI output—From the **Devices** > **Device Management** > **Cluster** > **General** area, you can view a set of pre-defined CLI outputs that can help you troubleshoot the cluster. The following commands are automatically run for the cluster:
  - **show running-config cluster**
  - **show cluster info**
  - **show cluster info health**
  - **show cluster info transport cp**
  - **show version**
  - **show asp drop**
  - **show counters**
  - **show arp**
  - **show int ip brief**
  - **show blocks**
  - **show cpu detailed**
  - **show interface** *ccl_interface*
  - **ping** *ccl_ip* **size** *ccl_mtu* **repeat 2**

  You can also enter any **show** command in the Command field. See View CLI Output, on page 132 for more information.

Perform a Ping on the Cluster Control Link

When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the ping fails, a notification is generated so you can fix

the MTU mismatch on connecting switches and try again. This tool lets you manually ping all nodes that have already joined the cluster in case you are having cluster control link connectivity problems.

Procedure

Step 1 Choose **Devices** > **Device Management**, click the **More** (⋮) icon next to the cluster, and choose **Cluster Live Status**.

Figure 297: Cluster Status

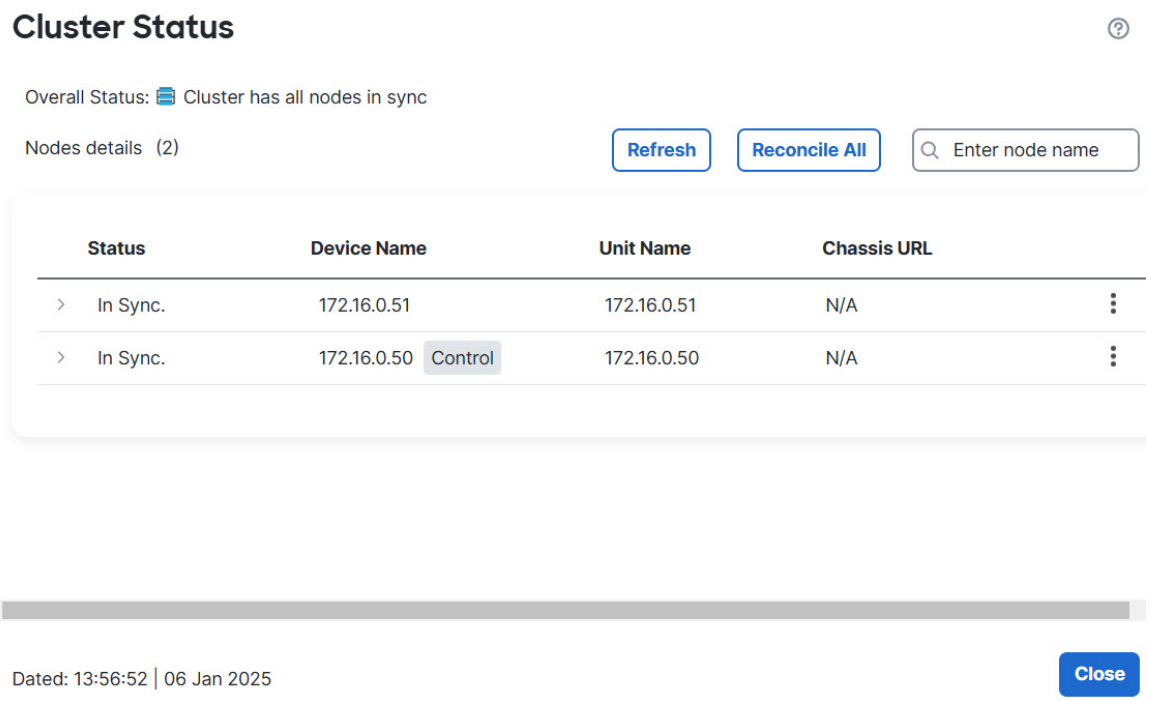

Step 2 Expand one of the nodes, and click **CCL Ping**.

Figure 298: CCL Ping

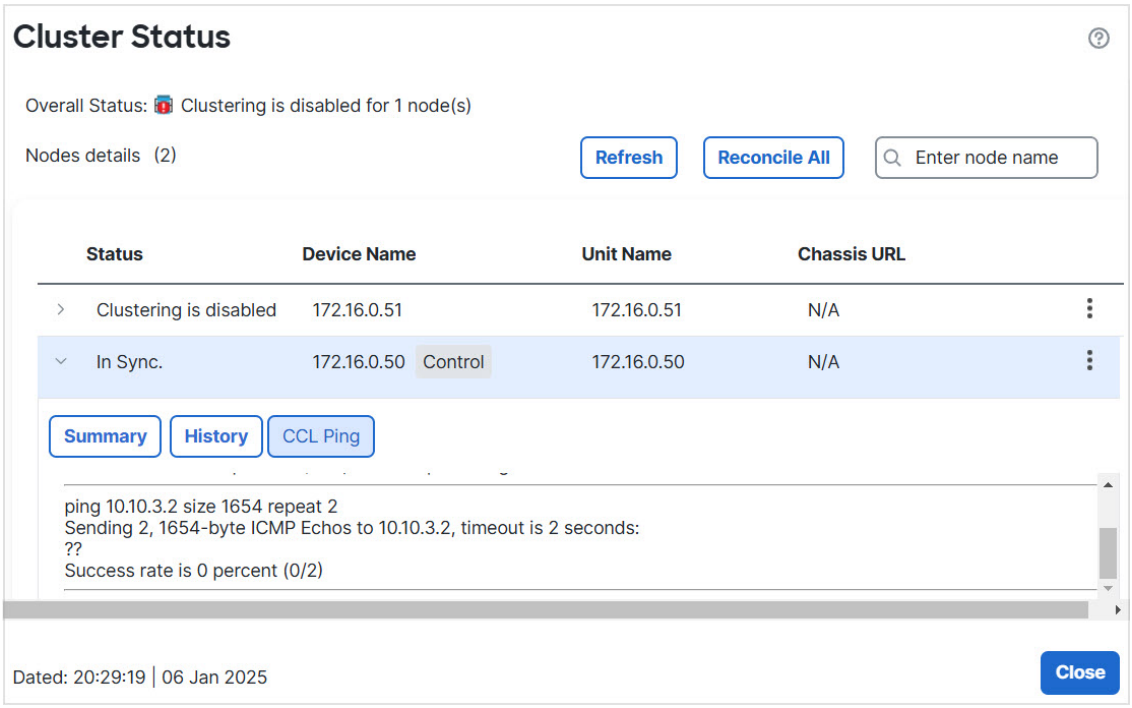

The node sends a ping on the cluster control link to every other node using a packet size that matches the maximum MTU.

Upgrading the Cluster

Perform the following steps to upgrade a Firewall Threat Defense Virtual cluster:

Before you begin

- Before you upgrade a cluster in the public cloud, copy the target version image to your cloud image repository and update the image ID in the cluster deployment template (we actually recommend replacing the existing template with a modified copy). This ensures that after the upgrade, new instances — for example, instances launched during cluster scaling — will use the correct version. If the marketplace does not have the image you need, such as when the cluster has been patched, create a custom image from a snapshot of a standalone Firewall Threat Defense Virtual instance running the correct version, with no instance-specific (day 0) configurations.
- For Firewall Threat Defense Virtual for AWS, suspend the HealthCheck and ReplaceUnhealthy processes before autoscaled cluster upgrade. This ensures that instances are not terminated by the Auto Scaling group during the post-upgrade reboot. You can resume the suspended processes afterwards. For instructions, see the Amazon EC2 Auto Scaling user guide: Suspend and resume Amazon EC2 Auto Scaling processes.

Procedure

Step 1 Upload the target image version to the cloud image storage.

Step 2 Update the cloud instance template of the cluster with the updated target image version.

a) Create a copy of the instance template with the target image version.

b) Attach the newly created template to cluster instance group.

Step 3 Upload the target image version upgrade package to the Firewall Management Center.

Step 4 Perform readiness check on the cluster that you want to upgrade.

Step 5 After successful readiness check, initiate installation of upgrade package.

Step 6 The Firewall Management Center upgrades the cluster nodes one at a time.

Step 7 The Firewall Management Center displays a notification after successful upgrade of the cluster.

There is no change in the serial number and UUID of the instance after the upgrade.

Reference for Clustering

This section includes more information about how clustering operates.

Threat Defense Features and Clustering

Some Firewall Threat Defense features are not supported with clustering, and some are only supported on the control unit. Other features might have caveats for proper usage.

Unsupported Features and Clustering

These features cannot be configured with clustering enabled, and the commands will be rejected.

Note To view FlexConfig features that are also not supported with clustering, for example WCCP inspection, see the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the Firewall Management Center GUI. See FlexConfig Policies, on page 2667.

- Remote access VPN (SSL VPN and IPsec VPN)
- DHCP client, server, and proxy. DHCP relay is supported.
- Virtual Tunnel Interfaces (VTIs)
- High Availability
- Integrated Routing and Bridging
- Firewall Management Center UCAPL/CC mode

Centralized Features for Clustering

The following features are only supported on the control node, and are not scaled for the cluster.

Note Traffic for centralized features is forwarded from member nodes to the control node over the cluster control link.

If you use the rebalancing feature, traffic for centralized features may be rebalanced to non-control nodes before the traffic is classified as a centralized feature; if this occurs, the traffic is then sent back to the control node.

For centralized features, if the control node fails, all connections are dropped, and you have to re-establish the connections on the new control node.

Note To view FlexConfig features that are also centralized with clustering, for example RADIUS inspection, see the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the Firewall Management Center GUI. See FlexConfig Policies, on page 2667.

- The following application inspections:
  - DCERPC
  - ESMTP
  - NetBIOS
  - PPTP
  - RSH
  - SQLNET
  - SUNRPC
  - TFTP
  - XDMCP
- Static route monitoring

Cisco Trustsec and Clustering

Only the control node learns security group tag (SGT) information. The control node then populates the SGT to data nodes, and data nodes can make a match decision for SGT based on the security policy.

Connection Settings and Clustering

Connection limits are enforced cluster-wide. Each node has an estimate of the cluster-wide counter values based on broadcast messages. Due to efficiency considerations, the configured connection limit across the cluster might not be enforced exactly at the limit number. Each node may overestimate or underestimate the

cluster-wide counter value at any given time. However, the information will get updated over time in a load-balanced cluster.

Dynamic Routing and Clustering

In Individual interface mode, each node runs the routing protocol as a standalone router, and routes are learned by each node independently.

Figure 299: Dynamic Routing in Individual Interface Mode

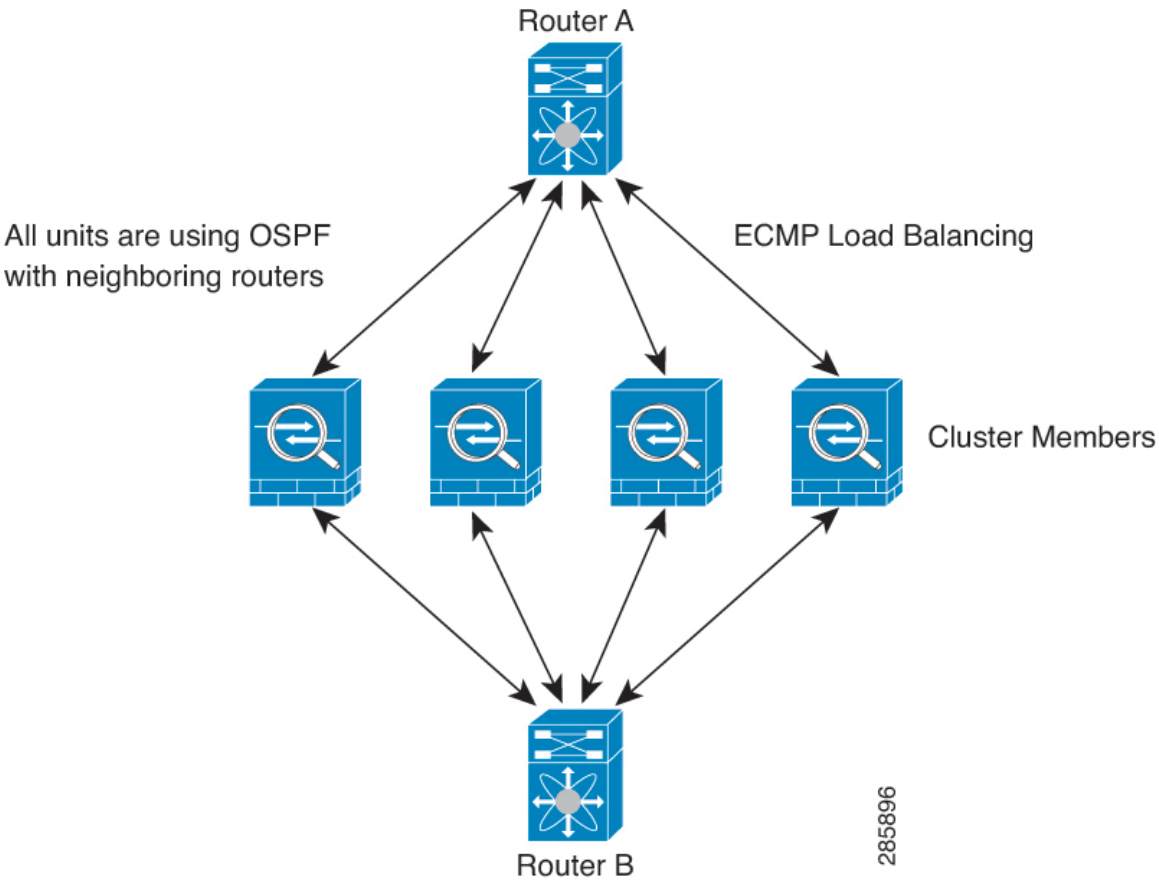

In the above diagram, Router A learns that there are 4 equal-cost paths to Router B, each through a node. ECMP is used to load balance traffic between the 4 paths. Each node picks a different router ID when talking to external routers.

You must configure a cluster pool for the router ID so that each node has a separate router ID.

FTP and Clustering

- If FTP data channel and control channel flows are owned by different cluster members, then the data channel owner will periodically send idle timeout updates to the control channel owner and update the idle timeout value. However, if the control flow owner is reloaded, and the control flow is re-hosted, the parent/child flow relationship will not longer be maintained; the control flow idle timeout will not be updated.

NAT and Clustering

For NAT usage, see the following limitations.

NAT can affect the overall throughput of the cluster. Inbound and outbound NAT packets can be sent to different Firewall Threat Defenses in the cluster, because the load balancing algorithm relies on IP addresses and ports, and NAT causes inbound and outbound packets to have different IP addresses and/or ports. When a packet arrives at the Firewall Threat Defense that is not the NAT owner, it is forwarded over the cluster control link to the owner, causing large amounts of traffic on the cluster control link. Note that the receiving node does not create a forwarding flow to the owner, because the NAT owner may not end up creating a connection for the packet depending on the results of security and policy checks.

If you still want to use NAT in clustering, then consider the following guidelines:

- No Proxy ARP—For Individual interfaces, a proxy ARP reply is never sent for mapped addresses. This prevents the adjacent router from maintaining a peer relationship with an ASA that may no longer be in the cluster. The upstream router needs a static route or PBR with Object Tracking for the mapped addresses that points to the Main cluster IP address.

- PAT with Port Block Allocation—See the following guidelines for this feature:
  - Maximum-per-host limit is not a cluster-wide limit, and is enforced on each node individually. Thus, in a 3-node cluster with the maximum-per-host limit configured as 1, if the traffic from a host is load-balanced across all 3 nodes, then it can get allocated 3 blocks with 1 in each node.
  - Port blocks created on the backup node from the backup pools are not accounted for when enforcing the maximum-per-host limit.
  - On-the-fly PAT rule modifications, where the PAT pool is modified with a completely new range of IP addresses, will result in xlate backup creation failures for the xlate backup requests that were still in transit while the new pool became effective. This behavior is not specific to the port block allocation feature, and is a transient PAT pool issue seen only in cluster deployments where the pool is distributed and traffic is load-balanced across the cluster nodes.
  - When operating in a cluster, you cannot simply change the block allocation size. The new size is effective only after you reload each device in the cluster. To avoid having to reload each device, we recommend that you delete all block allocation rules and clear all xlates related to those rules. You can then change the block size and recreate the block allocation rules.

- NAT pool address distribution for dynamic PAT—When you configure a PAT pool, the cluster divides each IP address in the pool into port blocks. By default, each block is 512 ports, but if you configure port block allocation rules, your block setting is used instead. These blocks are distributed evenly among the nodes in the cluster, so that each node has one or more blocks for each IP address in the PAT pool. Thus, you could have as few as one IP address in a PAT pool for a cluster, if that is sufficient for the number of PAT'ed connections you expect. Port blocks cover the 1024-65535 port range, unless you configure the option to include the reserved ports, 1-1023, on the PAT pool NAT rule.
- Reusing a PAT pool in multiple rules—To use the same PAT pool in multiple rules, you must be careful about the interface selection in the rules. You must either use specific interfaces in all rules, or "any" in all rules. You cannot mix specific interfaces and "any" across the rules, or the system might not be able to match return traffic to the right node in the cluster. Using unique PAT pools per rule is the most reliable option.
- No round-robin—Round-robin for a PAT pool is not supported with clustering.
- No extended PAT—Extended PAT is not supported with clustering.

- Dynamic NAT xlates managed by the control node—The control node maintains and replicates the xlate table to data nodes. When a data node receives a connection that requires dynamic NAT, and the xlate is not in the table, it requests the xlate from the control node. The data node owns the connection.
- Stale xlates—The xlate idle time on the connection owner does not get updated. Thus, the idle time might exceed the idle timeout. An idle timer value higher than the configured timeout with a refcnt of 0 is an indication of a stale xlate.
- No static PAT for the following inspections—
  - FTP
  - RSH
  - SQLNET
  - TFTP
  - XDMCP
  - SIP
- If you have an extremely large number of NAT rules, over ten thousand, you should enable the transactional commit model using the **asp rule-engine transactional-commit nat** command in the device CLI. Otherwise, the node might not be able to join the cluster.

SIP Inspection and Clustering

A control flow can be created on any node (due to load balancing); its child data flows must reside on the same node.

SNMP and Clustering

You should always use the Local address, and not the Main cluster IP address for SNMP polling. If the SNMP agent polls the Main cluster IP address, if a new control node is elected, the poll to the new control node will fail.

When using SNMPv3 with clustering, if you add a new cluster node after the initial cluster formation, then SNMPv3 users are not replicated to the new node. You must remove the users, and re-add them, and then redeploy your configuration to force the users to replicate to the new node.

Syslog and Clustering

- Each node in the cluster generates its own syslog messages. You can configure logging so that each node uses either the same or a different device ID in the syslog message header field. For example, the hostname configuration is replicated and shared by all nodes in the cluster. If you configure logging to use the hostname as the device ID, syslog messages generated by all nodes look as if they come from a single node. If you configure logging to use the local-node name that is assigned in the cluster bootstrap configuration as the device ID, syslog messages look as if they come from different nodes.

VPN and Clustering

Site-to-site VPN is a centralized feature; only the control node supports VPN connections.

Note Remote access VPN is not supported with clustering.

VPN functionality is limited to the control node and does not take advantage of the cluster high availability capabilities. If the control node fails, all existing VPN connections are lost, and VPN users will see a disruption in service. When a new control node is elected, you must reestablish the VPN connections.

For connections to an Individual interface when using PBR or ECMP, you must always connect to the Main cluster IP address, not a Local address.

VPN-related keys and certificates are replicated to all nodes.

Performance Scaling Factor

When you combine multiple units into a cluster, you can expect the total cluster performance to be approximately 80% of the maximum combined throughput.

For example, if your model can handle approximately 10 Gbps of traffic when running alone, then for a cluster of 8 units, the maximum combined throughput will be approximately 80% of 80 Gbps (8 units x 10 Gbps): 64 Gbps.

Control Node Election

Nodes of the cluster communicate over the cluster control link to elect a control node as follows:

1. When you enable clustering for a node (or when it first starts up with clustering already enabled), it broadcasts an election request every 3 seconds.
2. Any other nodes with a higher priority respond to the election request; the priority is set between 1 and 100, where 1 is the highest priority.
3. If after 45 seconds, a node does not receive a response from another node with a higher priority, then it becomes the control node.

Note If multiple nodes tie for the highest priority, the cluster node name and then the serial number is used to determine the control node.

4. If a node later joins the cluster with a higher priority, it does not automatically become the control node; the existing control node always remains as the control node unless it stops responding, at which point a new control node is elected.
5. In a "split brain" scenario when there are temporarily multiple control nodes, then the node with highest priority retains the role while the other nodes return to data node roles.

Note You can manually force a node to become the control node. For centralized features, if you force a control node change, then all connections are dropped, and you have to re-establish the connections on the new control node.

High Availability within the Cluster

Clustering provides high availability by monitoring node and interface health and by replicating connection states between nodes.

Node Health Monitoring

Each node periodically sends a broadcast heartbeat packet over the cluster control link. If the control node does not receive any heartbeat packets or other packets from a data node within the configurable timeout period, then the control node removes the data node from the cluster. If the data nodes do not receive packets from the control node, then a new control node is elected from the remaining nodes.

If nodes cannot reach each other over the cluster control link because of a network failure and not because a node has actually failed, then the cluster may go into a "split brain" scenario where isolated data nodes will elect their own control nodes. For example, if a router fails between two cluster locations, then the original control node at location 1 will remove the location 2 data nodes from the cluster. Meanwhile, the nodes at location 2 will elect their own control node and form their own cluster. Note that asymmetric traffic may fail in this scenario. After the cluster control link is restored, then the control node that has the higher priority will keep the control node's role.

Interface Monitoring

Each node monitors the link status of all named hardware interfaces in use, and reports status changes to the control node.

All physical interfaces are monitored; only named interfaces can be monitored. You can optionally disable monitoring per interface.

A node is removed from the cluster if its monitored interfaces fail. The node is removed after 500 ms.

Status After Failure

If the control node fails, then another member of the cluster with the highest priority (lowest number) becomes the control node.

The Firewall Threat Defense automatically tries to rejoin the cluster, depending on the failure event.

Note When the Firewall Threat Defense becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the Management interface can send and receive traffic.

Rejoining the Cluster

After a cluster member is removed from the cluster, how it can rejoin the cluster depends on why it was removed:

- Failed cluster control link when initially joining—After you resolve the problem with the cluster control link, you must manually rejoin the cluster by re-enabling clustering.
- Failed cluster control link after joining the cluster—The Firewall Threat Defense automatically tries to rejoin every 5 minutes, indefinitely.
- Failed data interface—The Firewall Threat Defense automatically tries to rejoin at 5 minutes, then at 10 minutes, and finally at 20 minutes. If the join is not successful after 20 minutes, then the Firewall Threat

Defense application disables clustering. After you resolve the problem with the data interface, you have to manually enable clustering.

- Failed node—If the node was removed from the cluster because of a node health check failure, then rejoining the cluster depends on the source of the failure. For example, a temporary power failure means the node will rejoin the cluster when it starts up again as long as the cluster control link is up. The Firewall Threat Defense application attempts to rejoin the cluster every 5 seconds.
- Internal error—Internal failures include: application sync timeout; inconsistent application statuses; and so on.
- Failed configuration deployment—If you deploy a new configuration from Firewall Management Center, and the deployment fails on some cluster members but succeeds on others, then the nodes that failed are removed from the cluster. You must manually rejoin the cluster by re-enabling clustering. If the deployment fails on the control node, then the deployment is rolled back, and no members are removed. If the deployment fails on all data nodes, then the deployment is rolled back, and no members are removed.

Data Path Connection State Replication

Every connection has one owner and at least one backup owner in the cluster. The backup owner does not take over the connection in the event of a failure; instead, it stores TCP/UDP state information, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner is usually also the director.

Some traffic requires state information above the TCP or UDP layer. See the following table for clustering support or lack of support for this kind of traffic.

Table 44: Features Replicated Across the Cluster

Traffic	State Support	Notes
Up time	Yes	Keeps track of the system up time.
ARP Table	Yes	—
MAC address table	Yes	—
User Identity	Yes	—
IPv6 Neighbor database	Yes	—
Dynamic routing	Yes	—
SNMP Engine ID	**No**	—

How the Cluster Manages Connections

Connections can be load-balanced to multiple nodes of the cluster. Connection roles determine how connections are handled in both normal operation and in a high availability situation.

Connection Roles

See the following roles defined for each connection:

- Owner—Usually, the node that initially receives the connection. The owner maintains the TCP state and processes packets. A connection has only one owner. If the original owner fails, then when new nodes receive packets from the connection, the director chooses a new owner from those nodes.
- Backup owner—The node that stores TCP/UDP state information received from the owner, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner does not take over the connection in the event of a failure. If the owner becomes unavailable, then the first node to receive packets from the connection (based on load balancing) contacts the backup owner for the relevant state information so it can become the new owner.

  As long as the director (see below) is not the same node as the owner, then the director is also the backup owner. If the owner chooses itself as the director, then a separate backup owner is chosen.

  For clustering on the Firepower 9300, which can include up to 3 cluster nodes in one chassis, if the backup owner is on the same chassis as the owner, then an additional backup owner will be chosen from another chassis to protect flows from a chassis failure.
- Director—The node that handles owner lookup requests from forwarders. When the owner receives a new connection, it chooses a director based on a hash of the source/destination IP address and ports (see below for ICMP hash details), and sends a message to the director to register the new connection. If packets arrive at any node other than the owner, the node queries the director about which node is the owner so it can forward the packets. A connection has only one director. If a director fails, the owner chooses a new director.

  As long as the director is not the same node as the owner, then the director is also the backup owner (see above). If the owner chooses itself as the director, then a separate backup owner is chosen.

  ICMP/ICMPv6 hash details:

  - For Echo packets, the source port is the ICMP identifier, and the destination port is 0.
  - For Reply packets, the source port is 0, and the destination port is the ICMP identifier.
  - For other packets, both source and destination ports are 0.
- Forwarder—A node that forwards packets to the owner. If a forwarder receives a packet for a connection it does not own, it queries the director for the owner, and then establishes a flow to the owner for any other packets it receives for this connection. The director can also be a forwarder. Note that if a forwarder receives the SYN-ACK packet, it can derive the owner directly from a SYN cookie in the packet, so it does not need to query the director. (If you disable TCP sequence randomization, the SYN cookie is not used; a query to the director is required.) For short-lived flows such as DNS and ICMP, instead of querying, the forwarder immediately sends the packet to the director, which then sends them to the owner. A connection can have multiple forwarders; the most efficient throughput is achieved by a good load-balancing method where there are no forwarders and all packets of a connection are received by the owner.

Note We do not recommend disabling TCP sequence randomization when using clustering. There is a small chance that some TCP sessions won't be established, because the SYN/ACK packet might be dropped.

- Fragment Owner—For fragmented packets, cluster nodes that receive a fragment determine a fragment owner using a hash of the fragment source IP address, destination IP address, and the packet ID. All fragments are then forwarded to the fragment owner over the cluster control link. Fragments may be

load-balanced to different cluster nodes, because only the first fragment includes the 5-tuple used in the switch load balance hash. Other fragments do not contain the source and destination ports and may be load-balanced to other cluster nodes. The fragment owner temporarily reassembles the packet so it can determine the director based on a hash of the source/destination IP address and ports. If it is a new connection, the fragment owner will register to be the connection owner. If it is an existing connection, the fragment owner forwards all fragments to the provided connection owner over the cluster control link. The connection owner will then reassemble all fragments.

Port Address Translation Connections

New Connection Ownership

Traffic redirection is not supported in this release. When a new connection is directed to a node of the cluster via load balancing, that node owns both directions of the connection. All the subsequent packets for the same connection should arrive the same node. If any connection packets arrive at a different node, they will be dropped. If a reverse flow arrives at a different node, it will be dropped as well. For centralized features, if the connections do not arrive on the control node, they will be dropped.

By default, AWS GWLB uses 5-tuple to maintain flow stickiness. It is recommended to enable 2-tuple or 3-tuple stickiness on AWS GWLB to ensure the same flows are sent to the same node.

Sample Data Flow for TCP

The following example shows the establishment of a new connection.

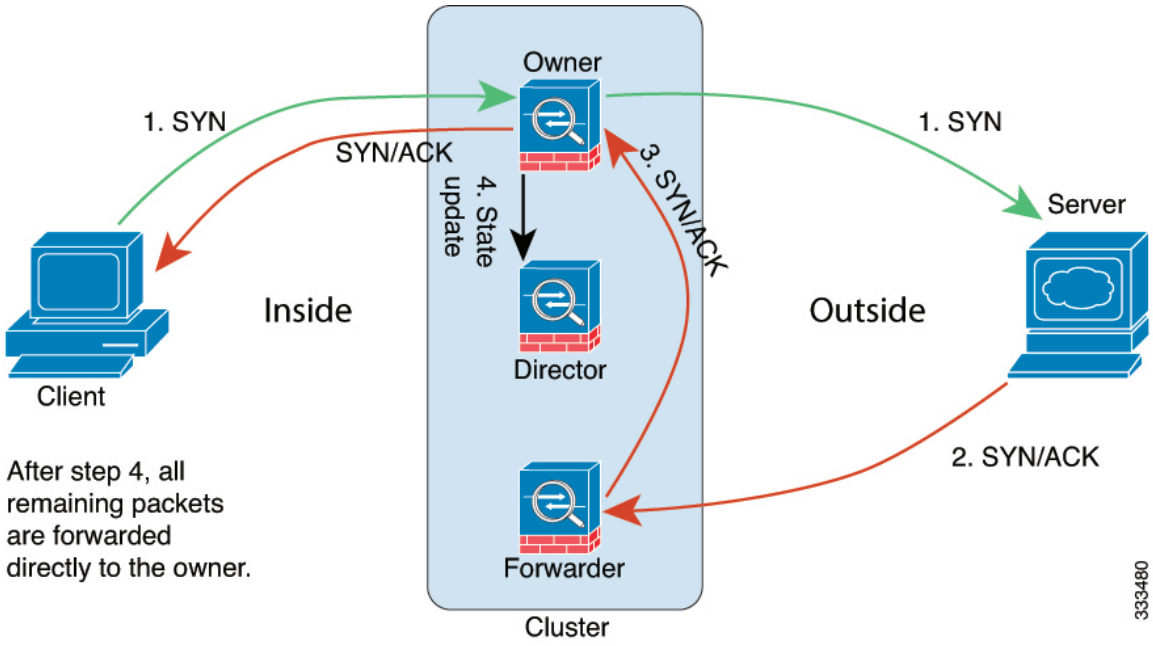

1. The SYN packet originates from the client and is delivered to one Firewall Threat Defense (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.
2. The SYN-ACK packet originates from the server and is delivered to a different Firewall Threat Defense (based on the load balancing method). This Firewall Threat Defense is the forwarder.

3. Because the forwarder does not own the connection, it decodes owner information from the SYN cookie, creates a forwarding flow to the owner, and forwards the SYN-ACK to the owner.
4. The owner sends a state update to the director, and forwards the SYN-ACK to the client.
5. The director receives the state update from the owner, creates a flow to the owner, and records the TCP state information as well as the owner. The director acts as the backup owner for the connection.
6. Any subsequent packets delivered to the forwarder will be forwarded to the owner.
7. If packets are delivered to any additional nodes, it will query the director for the owner and establish a flow.
8. Any state change for the flow results in a state update from the owner to the director.

Sample Data Flow for ICMP and UDP

The following example shows the establishment of a new connection.

1. ***Figure 300: ICMP and UDP Data Flow***

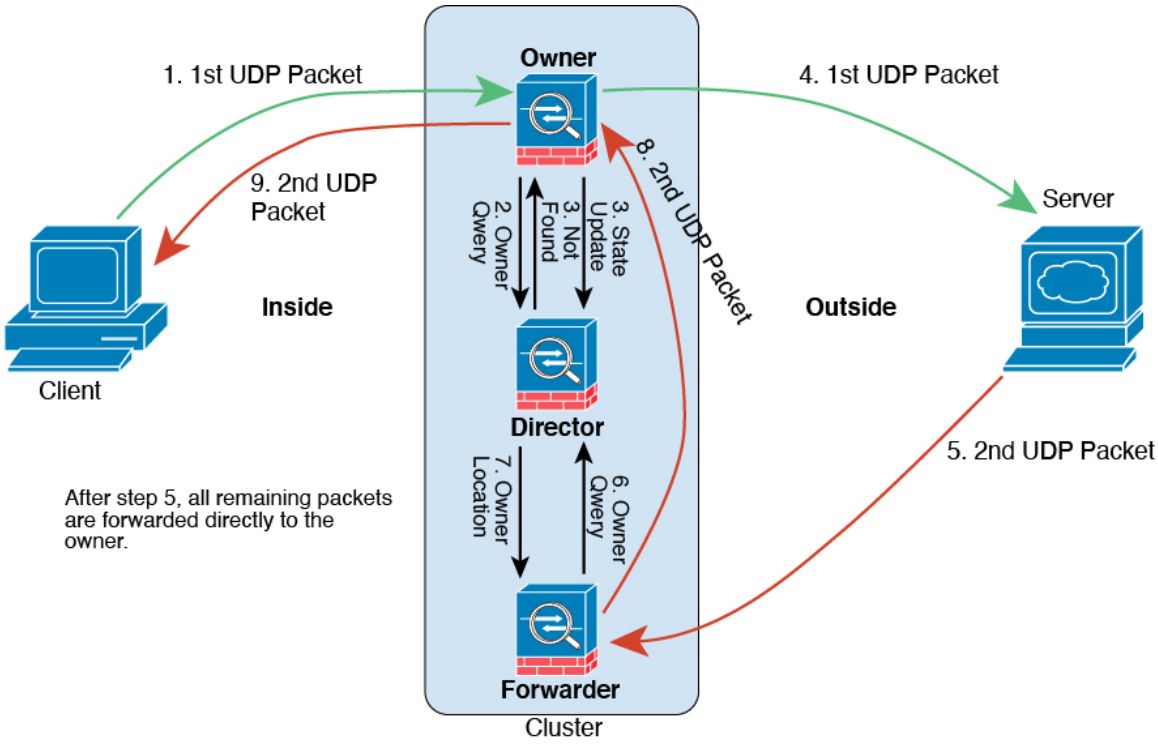

The first UDP packet originates from the client and is delivered to one Firewall Threat Defense (based on the load balancing method).

2. The node that received the first packet queries the director node that is chosen based on a hash of the source/destination IP address and ports.
3. The director finds no existing flow, creates a director flow and forwards the packet back to the previous node. In other words, the director has elected an owner for this flow.
4. The owner creates the flow, sends a state update to the director, and forwards the packet to the server.

5. The second UDP packet originates from the server and is delivered to the forwarder.
6. The forwarder queries the director for ownership information. For short-lived flows such as DNS, instead of querying, the forwarder immediately sends the packet to the director, which then sends it to the owner.
7. The director replies to the forwarder with ownership information.
8. The forwarder creates a forwarding flow to record owner information and forwards the packet to the owner.
9. The owner forwards the packet to the client.

History for Threat Defense Virtual Clustering in the Public Cloud

Table 45:

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
MTU ping test on cluster node join	7.6.0	7.6.0	When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the ping fails, a notification is generated so you can fix the MTU mismatch on connecting switches and try again.
Cluster control link ping tool.	7.2.6/7.4.1	Any	You can check to make sure all the cluster nodes can reach each other over the cluster control link by performing a ping. One major cause for the failure of a node to join the cluster is an incorrect cluster control link configuration; for example, the cluster control link MTU may be set higher than the connecting switch MTUs. New/modified screens: **Devices** > **Device Management** > **More** > **Cluster Live Status.** Version restrictions: Not supported with Firewall Management Center Version 7.3.x or 7.4.0.
Troubleshooting file generation and download available from Device and Cluster pages.	7.4.1	7.4.1	You can generate and download troubleshooting files for each device on the Device page and also for all cluster nodes on the Cluster page. For a cluster, you can download all files as a single compressed file. You can also include cluster logs for the cluster for cluster nodes. You can alternatively trigger file generation from the **Devices** > **Device Management** > **More** > **Troubleshoot Files** menu. New/modified screens: • **Devices** > **Device Management** > **Device** > **General** • **Devices** > **Device Management** > **Cluster** > **General**

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
View CLI output for a device or device cluster.	7.4.1	Any	You can view a set of pre-defined CLI outputs that can help you troubleshoot the device or cluster. You can also enter any **show** command and see the output. New/modified screens: **Devices** > **Device Management** > **Cluster** > **General**
Cluster health monitor settings.	7.3.0	Any	You can now edit cluster health monitor settings. New/Modified screens: **Devices** > **Device Management** > **Cluster** > **Cluster Health Monitor Settings** **Note** If you previously configured these settings using FlexConfig, be sure to remove the FlexConfig configuration before you deploy. Otherwise the FlexConfig configuration will overwrite the management center configuration.
Cluster health monitor dashboard.	7.3.0	Any	You can now view cluster health on the cluster health monitor dashboard. New/Modified screens: **System** > **Health** > **Monitor**
Clustering for the Firewall Threat Defense Virtual in Azure.	7.3.0	7.3.0	You can now configure clustering for up to 16 nodes the Firewall Threat Defense Virtual in Azure for the Azure Gateway Load Balancer or for external load balancers. New/modified screens: • **Devices** > **Device Management** > **Add Cluster** • **Devices** > **Device Management** > **More** menu • **Devices** > **Device Management** > **Cluster** Supported platforms: Firewall Threat Defense Virtual in Azure
Clustering for the Firewall Threat Defense Virtual in the Public Cloud (Amazon Web Services and Google Cloud Platform).	7.2.0	7.2.0	The Firewall Threat Defense Virtual supports Individual interface clustering for up to 16 nodes in the public cloud (AWS and GCP). New/Modified screens: • **Devices** > **Device Management** > **Add Device** • **Devices** > **Device Management** > **More** menu • **Devices** > **Device Management** > **Cluster** Supported platforms: Firewall Threat Defense Virtual in AWS and GCP

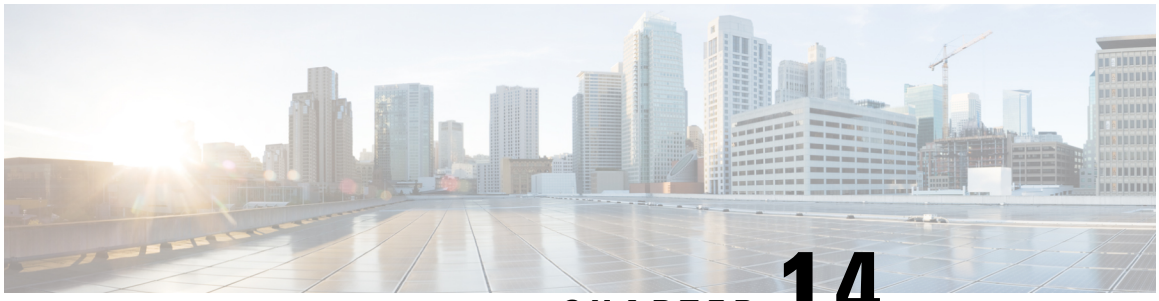

CHAPTER 14

Clustering for the Firepower 4100/9300

Clustering lets you group multiple Firewall Threat Defense nodes together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices.

Note Some features are not supported when using clustering. See Unsupported Features with Clustering, on page 753.

- About Clustering on the Firepower 4100/9300 Chassis, on page 699
- Licenses for Clustering, on page 703
- Requirements and Prerequisites for Clustering, on page 704
- Clustering Guidelines and Limitations, on page 707
- Configure Clustering, on page 711
- FXOS: Remove a Cluster Node, on page 737
- Firewall Management Center: Manage Cluster Members, on page 739
- Firewall Management Center: Monitoring the Cluster, on page 744
- Firewall Management Center: Troubleshooting the Cluster, on page 749
- Examples for Clustering, on page 751
- Reference for Clustering, on page 753
- History for Clustering, on page 765

About Clustering on the Firepower 4100/9300 Chassis

When you deploy a cluster on the Firepower 4100/9300 chassis, it does the following:

- For native instance clustering: Creates a *cluster-control link* (by default, port-channel 48) for node-to-node communication.

  For multi-instance clustering: You should pre-configure subinterfaces on one or more cluster-type EtherChannels; each instance needs its own cluster control link.

  For a cluster isolated to security modules within one Firepower 9300 chassis, this link utilizes the Firepower 9300 backplane for cluster communications.

  For clustering with multiple chassis, you need to manually assign physical interface(s) to this EtherChannel for communications between chassis.

- Creates the cluster bootstrap configuration within the application.

  When you deploy the cluster, the chassis supervisor pushes a minimal bootstrap configuration to each unit that includes the cluster name, cluster control link interface, and other cluster settings.

- Assigns data interfaces to the cluster as *Spanned* interfaces.

  For a cluster isolated to security modules within one Firepower 9300 chassis, spanned interfaces are not limited to EtherChannels, like it is for clustering with multiple chassis. The Firepower 9300 supervisor uses EtherChannel technology internally to load-balance traffic to multiple modules on a shared interface, so any data interface type works for Spanned mode. For clustering with multiple chassis, you must use Spanned EtherChannels for all data interfaces.

Note Individual interfaces are not supported, with the exception of a management interface.

- Assigns a management interface to all units in the cluster.

See the following sections for more information about clustering.

Bootstrap Configuration

When you deploy the cluster, the Firepower 4100/9300 chassis supervisor pushes a minimal bootstrap configuration to each unit that includes the cluster name, cluster control link interface, and other cluster settings.

Cluster Members

Cluster members work together to accomplish the sharing of the security policy and traffic flows.

One member of the cluster is the **control** unit. The control unit is determined automatically. All other members are **data** units.

You must perform all configuration on the control unit only; the configuration is then replicated to the data units.

Some features do not scale in a cluster, and the control unit handles all traffic for those features. .

Cluster Control Link

For native instance clustering: The cluster control link is automatically created using the Port-channel 48 interface.

For multi-instance clustering: You should pre-configure subinterfaces on one or more cluster-type EtherChannels; each instance needs its own cluster control link.

For a cluster isolated to security modules within one Firepower 9300 chassis, this interface has no member interfaces. This Cluster type EtherChannel utilizes the Firepower 9300 backplane for cluster communications. For clustering with multiple chassis, you must add one or more interfaces to the EtherChannel.

For a cluster with two chassis, do not directly connect the cluster control link from one chassis to the other chassis. If you directly connect the interfaces, then when one unit fails, the cluster control link fails, and thus

the remaining healthy unit fails. If you connect the cluster control link through a switch, then the cluster control link remains up for the healthy unit.

Cluster control link traffic includes both control and data traffic.

Size the Cluster Control Link

If possible, you should size the cluster control link to match the expected throughput of each chassis so the cluster control link can handle the worst-case scenarios.

Cluster control link traffic is comprised mainly of state update and forwarded packets. The amount of traffic at any given time on the cluster control link varies. The amount of forwarded traffic depends on the load-balancing efficacy or whether there is a lot of traffic for centralized features. For example:

- NAT results in poor load balancing of connections, and the need to rebalance all returning traffic to the correct units.
- When membership changes, the cluster needs to rebalance a large number of connections, thus temporarily using a large amount of cluster control link bandwidth.

A higher-bandwidth cluster control link helps the cluster to converge faster when there are membership changes and prevents throughput bottlenecks.

Note If your cluster has large amounts of asymmetric (rebalanced) traffic, then you should increase the cluster control link size.

Cluster Control Link Redundancy

The following diagram shows how to use an EtherChannel as a cluster control link in a Virtual Switching System (VSS), Virtual Port Channel (vPC), StackWise, or StackWise Virtual environment. All links in the EtherChannel are active. When the switch is part of a redundant system, then you can connect firewall interfaces within the same EtherChannel to separate switches in the redundant system. The switch interfaces are members of the same EtherChannel port-channel interface, because the separate switches act like a single switch. Note that this EtherChannel is device-local, not a Spanned EtherChannel.

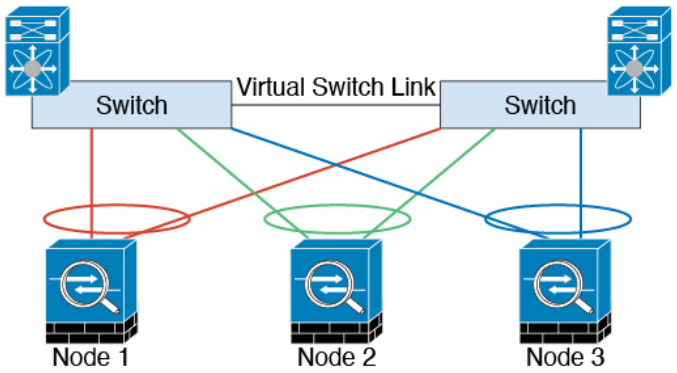

Cluster Control Link Reliability for Inter-Chassis Clustering

To ensure cluster control link functionality, be sure the round-trip time (RTT) between units is less than 20 ms. This maximum latency enhances compatibility with cluster members installed at different geographical sites. To check your latency, perform a ping on the cluster control link between units.

The cluster control link must be reliable, with no out-of-order or dropped packets; for example, for inter-site deployment, you should use a dedicated link.

Cluster Control Link Network

The Firepower 4100/9300 chassis auto-generates the cluster control link interface IP address for each unit based on the chassis ID and slot ID: 127.2.*chassis_id*.*slot_id*. For multi-instance clusters, which typically use different VLAN subinterfaces of the same EtherChannel, the same IP address can be used for different clusters because of VLAN separation. The cluster control link network cannot include any routers between units; only Layer 2 switching is allowed.

Management Network

We recommend connecting all units to a single management network. This network is separate from the cluster control link.

Management Interface

You must assign a Management type interface to the cluster. This interface is a special individual interface as opposed to a Spanned interface. The management interface lets you connect directly to each unit. This Management logical interface is separate from the other interfaces on the device. It is used to set up and register the device to the Secure Firewall Management Center. It uses its own local authentication, IP address, and static routing. Each cluster member uses a separate IP address on the management network that you set as part of the bootstrap configuration.

Cluster Interfaces

For a cluster isolated to security modules within one Firepower 9300 chassis, you can assign both physical interfaces or EtherChannels (also known as port channels) to the cluster. Interfaces assigned to the cluster are Spanned interfaces that load-balance traffic across all members of the cluster.

For clustering with multiple chassis, you can only assign data EtherChannels to the cluster. These Spanned EtherChannels include the same member interfaces on each chassis; on the upstream switch, all of these interfaces are included in a single EtherChannel, so the switch does not know that it is connected to multiple devices.

You can use regular firewall interfaces or IPS-only interfaces (inline sets or passive interfaces).

Individual interfaces are not supported, with the exception of a management interface.

Spanned EtherChannels

You can group one or more interfaces per chassis into an EtherChannel that spans all chassis in the cluster. The EtherChannel aggregates the traffic across all the available active interfaces in the channel.

For regular firewall interfaces:A Spanned EtherChannel can be configured in both routed and transparent firewall modes. In routed mode, the EtherChannel is configured as a routed interface with a single IP address. In transparent mode, the IP address is assigned to the BVI, not to the bridge group member interface.

The EtherChannel inherently provides load balancing as part of basic operation.

For multi-instance clusters, each cluster requires dedicated data EtherChannels; you cannot use shared interfaces or VLAN subinterfaces.

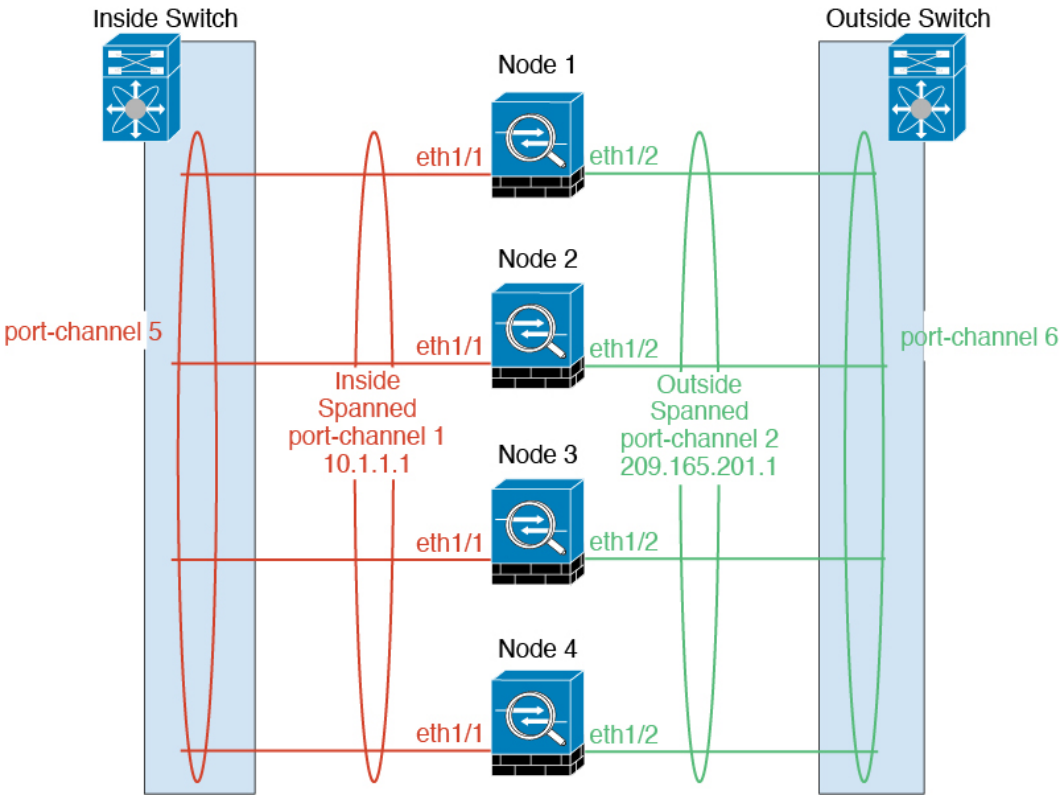

Connecting to a Redundant Switch System

We recommend connecting EtherChannels to a redundant switch system such as a VSS, vPC, StackWise, or StackWise Virtual system to provide redundancy for your interfaces.

Configuration Replication

All nodes in the cluster share a single configuration. You can only make configuration changes on the control node (with the exception of the bootstrap configuration), and changes are automatically synced to all other nodes in the cluster.

Licenses for Clustering

You assign feature licenses to the cluster as a whole, not to individual nodes. However, each node of the cluster consumes a separate license for each feature. The clustering feature itself does not require any licenses.

When you add a cluster node to the Firewall Management Center, you can specify the feature licenses you want to use for the cluster. You can modify licenses for the cluster in the **Devices** > **Device Management** > **Cluster** > **License** area.

Note If you add the cluster before the Firewall Management Center is licensed (and running in Evaluation mode), then when you license the Firewall Management Center, you can experience traffic disruption when you deploy policy changes to the cluster. Changing to licensed mode causes all data units to leave the cluster and then rejoin.

Requirements and Prerequisites for Clustering

Cluster Model Support

The Threat Defense supports clustering on the following models:

- Firepower 9300— You can include up to 16 nodes in the cluster. For example, you can use 1 module in 16 chassis, or 2 modules in 8 chassis, or any combination that provides a maximum of 16 modules. Supports clustering with multiple chassis and clustering isolated to security modules within one chassis.
- Firepower 4100—Supported for up to 16 nodes using clustering with multiple chassis.

User Roles

- Admin
- Access Admin
- Network Admin

Clustering Hardware and Software Requirements

All chassis in a cluster:

- Native instance clustering—For the Firepower 4100: All chassis must be the same model. For the Firepower 9300: All security modules must be the same type. For example, if you use clustering, all modules in the Firepower 9300 must be SM-40s. You can have different quantities of installed security modules in each chassis, although all modules present in the chassis must belong to the cluster including any empty slots.
- Container instance clustering—We recommend that you use the same security module or chassis model for each cluster instance. However, you can mix and match container instances on different Firepower 9300 security module types or Firepower 4100 models in the same cluster if required. You cannot mix Firepower 9300 and 4100 instances in the same cluster. For example, you can create a cluster using an instance on a Firepower 9300 SM-56, SM-48, and SM-40. Or you can create a cluster on a Firepower 4145 and a 4125.

- Must run the identical FXOS and application software except at the time of an image upgrade. Mismatched software versions can lead to poor performance, so be sure to upgrade all nodes in the same maintenance window.
- Must include the same interface configuration for interfaces you assign to the cluster, such as the same Management interface, EtherChannels, active interfaces, speed and duplex, and so on. You can use different network module types on the chassis as long as the capacity matches for the same interface IDs and interfaces can successfully bundle in the same spanned EtherChannel. Note that all data interfaces must be EtherChannels in clusters with multiple chassis. If you change the interfaces in FXOS after you enable clustering (by adding or removing interface modules, or configuring EtherChannels, for example), then perform the same changes on each chassis, starting with the data nodes, and ending with the control node.
- Must use the same NTP server. For Firewall Threat Defense, the Firewall Management Center must also use the same NTP server. Do not set the time manually.

Multi-Instance Clustering Requirements

- No intra-security-module/engine clustering—For a given cluster, you can only use a single container instance per security module/engine. You cannot add 2 container instances to the same cluster if they are running on the same module.

- Mix and match clusters and standalone instances—Not all container instances on a security module/engine need to belong to a cluster. You can use some instances as standalone or High Availability nodes. You can also create multiple clusters using separate instances on the same security module/engine.

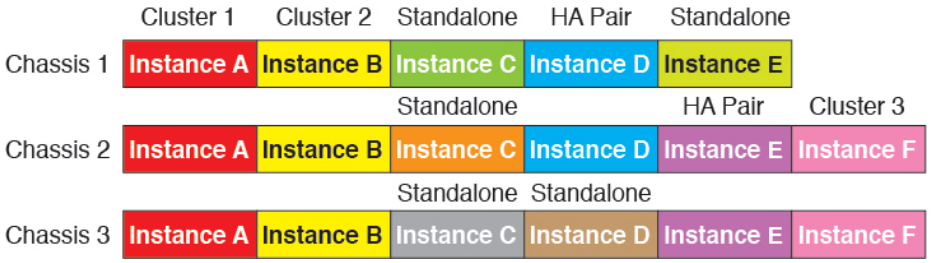

- All 3 modules in a Firepower 9300 must belong to the cluster—For the Firepower 9300, a cluster requires a single container instance on all 3 modules. You cannot create a cluster using instances on module 1 and 2, and then use a native instance on module 3, or example.

- Match resource profiles—We recommend that each node in the cluster use the same resource profile attributes; however, mismatched resources are allowed when changing cluster nodes to a different resource profile, or when using different models.

- Dedicated cluster control link—For clusters with multiple chassis, each cluster needs a dedicated cluster control link. For example, each cluster can use a separate subinterface on the same cluster-type EtherChannel, or use separate EtherChannels.

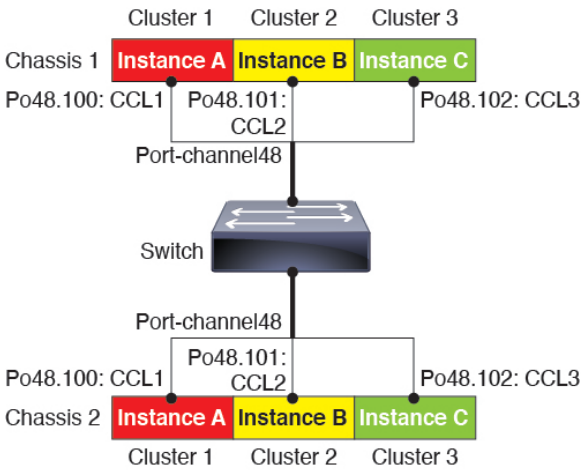

- No shared interfaces—Shared-type interfaces are not supported with clustering. However, the same Management and Eventing interfaces can used by multiple clusters.

- No subinterfaces—A multi-instance cluster cannot use FXOS-defined VLAN subinterfaces. An exception is made for the cluster control link, which can use a subinterface of the Cluster EtherChannel.

- Mix chassis models—We recommend that you use the same security module or chassis model for each cluster instance. However, you can mix and match container instances on different Firepower 9300

security module types or Firepower 4100 models in the same cluster if required. You cannot mix Firepower 9300 and 4100 instances in the same cluster. For example, you can create a cluster using an instance on a Firepower 9300 SM-56, SM-48, and SM-40. Or you can create a cluster on a Firepower 4145 and a 4125.

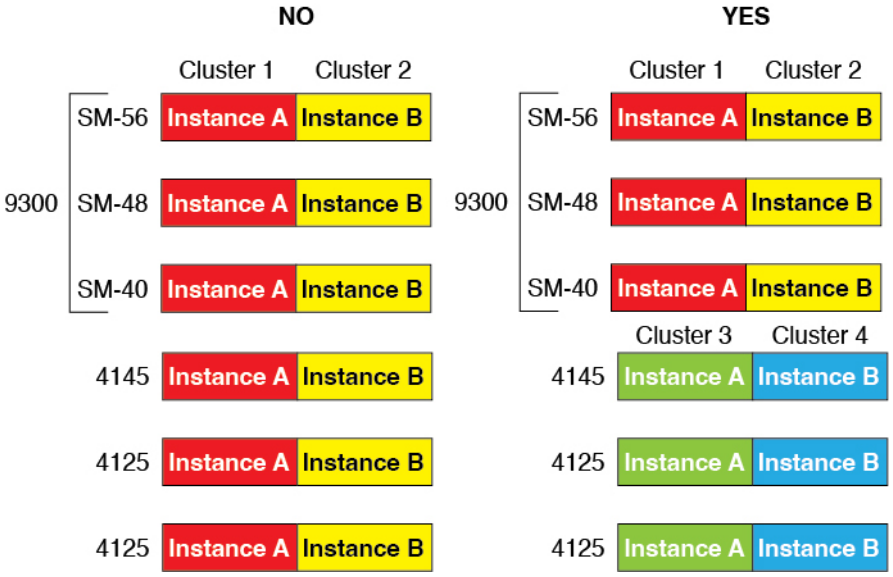

- Maximum 6 nodes—You can use up to six container instances in a cluster.

Switch Requirements

- Be sure to complete the switch configuration and successfully connect all the EtherChannels from the chassis to the switch(es) before you configure clustering on the Firepower 4100/9300 chassis.
- For supported switch characteristics, see Cisco FXOS Compatibility.

Clustering Guidelines and Limitations

Switches for Clustering

- Make sure connected switches match the MTU for both cluster data interfaces and the cluster control link interface. You should configure the cluster control link interface MTU to be at least 100 bytes higher than the data interface MTU, so make sure to configure the cluster control link connecting switch appropriately. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead. In addition, we do not recommend setting the cluster control link MTU between 2561 and 8362; due to block pool handling, this MTU size is not optimal for system operation. When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the ping fails, a notification is generated so you can fix the MTU mismatch on connecting switches and try again.
- For Cisco IOS XR systems, if you want to set a non-default MTU, set the IOS XR interface MTU to be 14 bytes higher than the cluster device MTU. Otherwise, OSPF adjacency peering attempts may fail

unless the **mtu-ignore** option is used. Note that the cluster device MTU should match the IOS XR *IPv4* MTU. This adjustment is not required for Cisco Catalyst and Cisco Nexus switches.

- On the switch(es) for the cluster control link interfaces, you can optionally enable Spanning Tree PortFast on the switch ports connected to the cluster unit to speed up the join process for new units.
- On the switch, we recommend that you use one of the following EtherChannel load-balancing algorithms: **source-dest-ip** or **src-dst-mixed-ip-port** (see the Cisco Nexus OS and Cisco IOS-XE **port-channel load-balance** command). Do not use a **vlan** keyword in the load-balance algorithm because it can cause unevenly distributed traffic to the devices in a cluster.
- If you change the load-balancing algorithm of the EtherChannel on the switch, the EtherChannel interface on the switch temporarily stops forwarding traffic, and the Spanning Tree Protocol restarts. There will be a delay before traffic starts flowing again.
- Switches on the cluster control link path should not verify the L4 checksum. Redirected traffic over the cluster control link does not have a correct L4 checksum. Switches that verify the L4 checksum could cause traffic to be dropped.
- Port-channel bundling downtime should not exceed the configured keepalive interval.
- On Supervisor 2T EtherChannels, the default hash distribution algorithm is adaptive. To avoid asymmetric traffic in a VSS design, change the hash algorithm on the port-channel connected to the cluster device to fixed:

  router(config)# **port-channel** *id* **hash-distribution fixed**

  Do not change the algorithm globally; you may want to take advantage of the adaptive algorithm for the VSS peer link.

- Firepower 4100/9300 clusters support LACP graceful convergence. So you can leave LACP graceful convergence enabled on connected Cisco Nexus switches.
- When you see slow bundling of a Spanned EtherChannel on the switch, you can enable LACP rate fast for an individual interface on the switch. FXOS EtherChannels have the LACP rate set to fast by default. Note that some switches, such as the Nexus series, do not support LACP rate fast when performing in-service software upgrades (ISSUs), so we do not recommend using ISSUs with clustering.

EtherChannels for Clustering

- In Catalyst 3750-X Cisco IOS software versions earlier than 15.1(1)S2, the cluster unit did not support connecting an EtherChannel to a switch stack. With default switch settings, if the cluster unit EtherChannel is connected cross stack, and if the control unit switch is powered down, then the EtherChannel connected to the remaining switch will not come up. To improve compatibility, set the **stack-mac persistent timer** command to a large enough value to account for reload time; for example, 8 minutes or 0 for indefinite. Or, you can upgrade to more a more stable switch software version, such as 15.1(1)S2.
- Spanned vs. Device-Local EtherChannel Configuration—Be sure to configure the switch appropriately for Spanned EtherChannels vs. Device-local EtherChannels.
  - Spanned EtherChannels—For cluster unit *Spanned* EtherChannels, which span across all members of the cluster, the interfaces are combined into a single EtherChannel on the switch. Make sure each interface is in the same channel group on the switch.

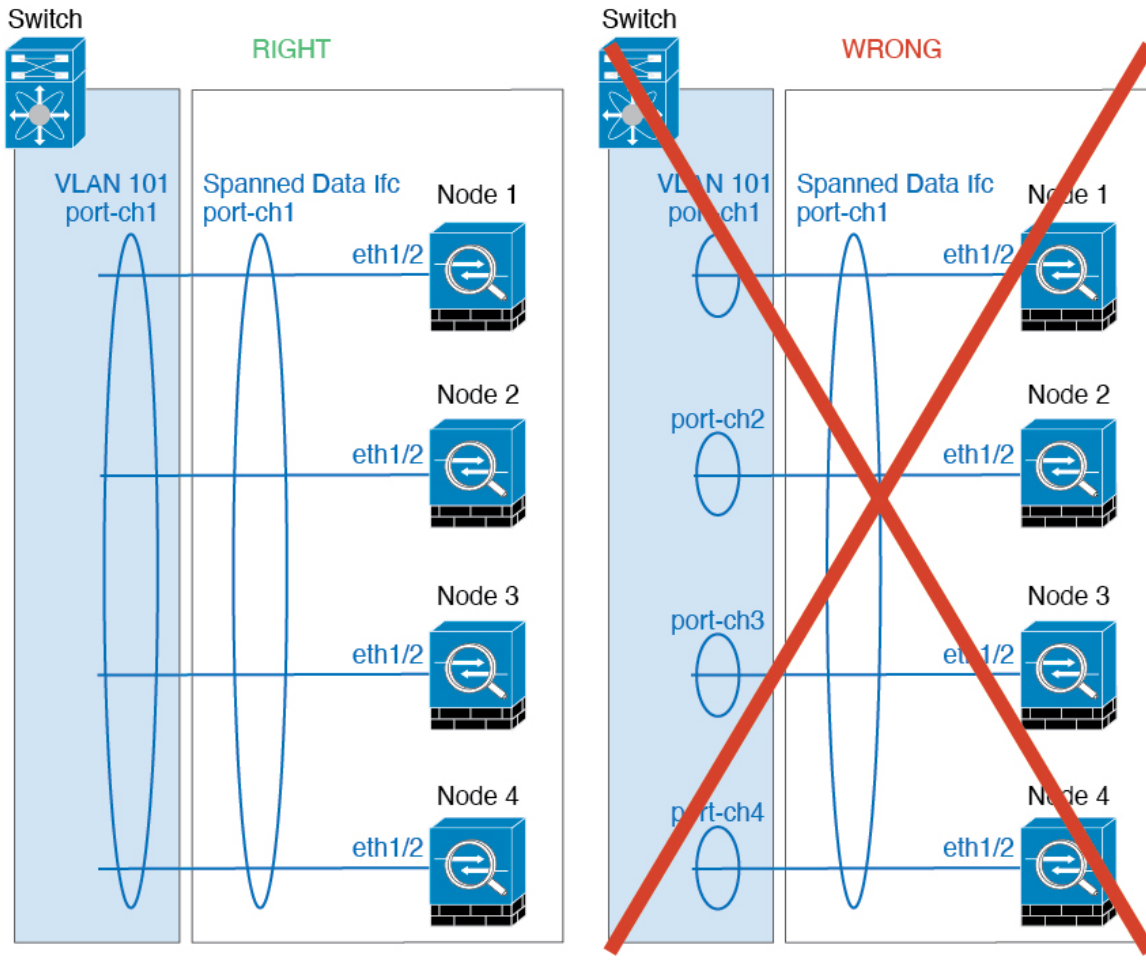

- Device-local EtherChannels—For cluster unit *Device-local* EtherChannels including any EtherChannels configured for the cluster control link, be sure to configure discrete EtherChannels on the switch; do not combine multiple cluster unit EtherChannels into one EtherChannel on the switch.

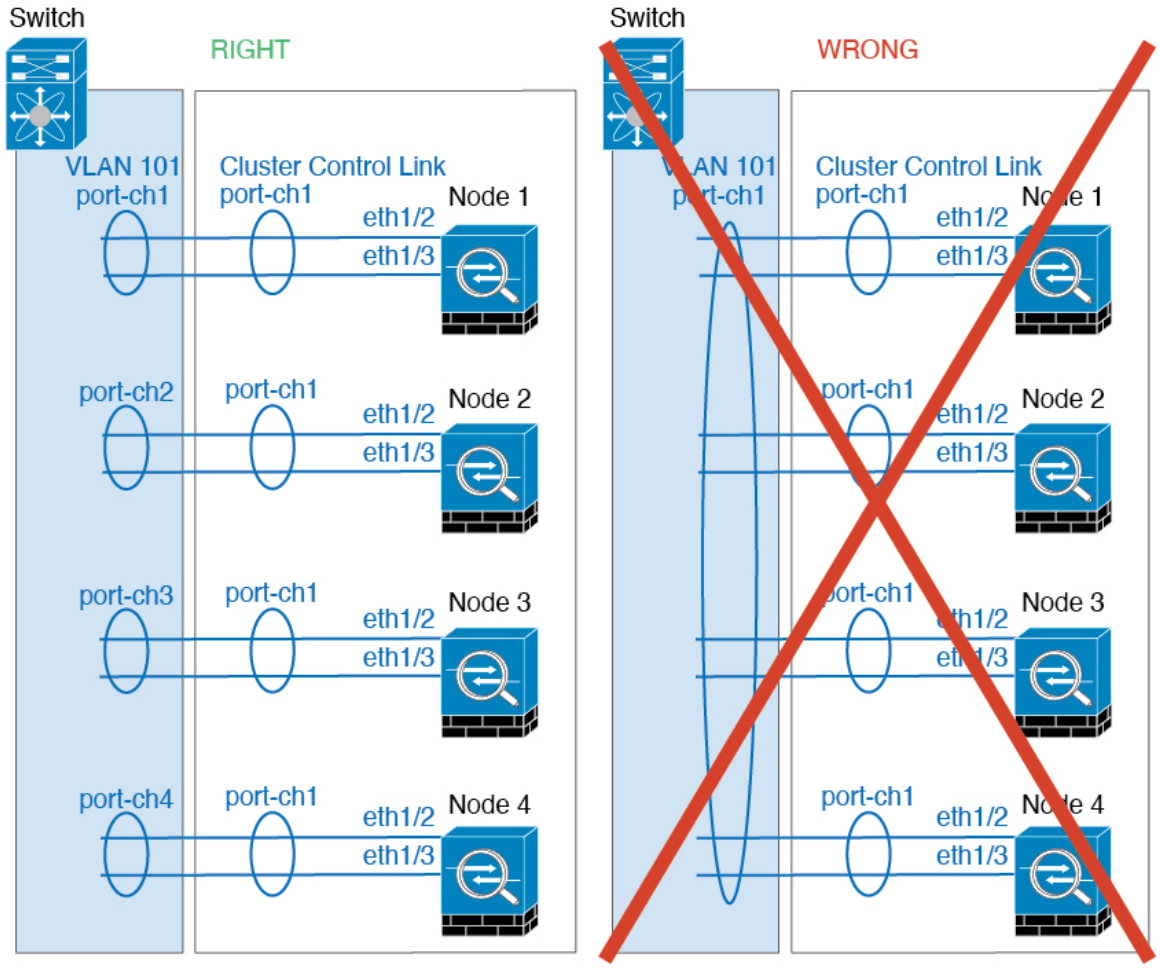

Additional Guidelines

- When significant topology changes occur (such as adding or removing an EtherChannel interface, enabling or disabling an interface on the Firepower 4100/9300 chassis or the switch, adding an additional switch to form a VSS, vPC, StackWise, or StackWise Virtual) you should disable the health check feature, and also disable interface monitoring for the disabled interfaces . When the topology change is complete, and the configuration change is synced to all units, you can re-enable the health check feature.
- When adding a unit to an existing cluster, or when reloading a unit, there will be a temporary, limited packet/connection drop; this is expected behavior. In some cases, the dropped packets can hang connections; for example, dropping a FIN/ACK packet for an FTP connection will make the FTP client hang. In this case, you need to reestablish the FTP connection.
- If you use a Windows 2003 server connected to a Spanned EtherChannel interface, when the syslog server port is down, and the server does not throttle ICMP error messages, then large numbers of ICMP messages are sent back to the cluster. These messages can result in some units of the cluster experiencing high CPU, which can affect performance. We recommend that you throttle ICMP error messages.
- We recommend connecting EtherChannels to a VSS, vPC, StackWise, or StackWise Virtual for redundancy.

- Within a chassis, you cannot cluster some security modules and run other security modules in standalone mode; you must include all security modules in the cluster.
- For decrypted TLS/SSL connections, the decryption states are not synchronized, and if the connection owner fails, then decrypted connections will be reset. New connections will need to be established to a new unit. Connections that are not decrypted (they match a do-not-decrypt rule) are not affected and are replicated correctly.

Defaults

- The cluster health check feature is enabled by default with the holdtime of 3 seconds. Interface health monitoring is enabled on all interfaces by default.
- The cluster auto-rejoin feature for a failed cluster control link is set to unlimited attempts every 5 minutes.
- The cluster auto-rejoin feature for a failed data interface is set to 3 attempts every 5 minutes, with the increasing interval set to 2.
- Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

Configure Clustering

You can easily deploy the cluster from the Firepower 4100/9300 supervisor. All initial configuration is automatically generated for each unit. You can then add the units to the Firewall Management Center and group them into a cluster.

FXOS: Add a Firewall Threat Defense Cluster

In native mode: You can add a cluster to a single Firepower 9300 chassis that is isolated to security modules within the chassis, or you can use multiple chassis.

In multi-instance mode: You can add one or more clusters to a single Firepower 9300 chassis that are isolated to security modules within the chassis (you must include an instance on each module), or add one or more clusters on multiple chassis.

For clusters on multiple chassis, you must configure each chassis separately. Add the cluster on one chassis; you can then copy the bootstrap configuration from the first chassis to the next chassis for ease of deployment

Create a Firewall Threat Defense Cluster

You can easily deploy the cluster from the Firepower 4100/9300 chassis supervisor. All initial configuration is automatically generated for each unit.

For clustering on multiple chassis, you must configure each chassis separately. Deploy the cluster on one chassis; you can then copy the bootstrap configuration from the first chassis to the next chassis for ease of deployment.

In a Firepower 9300 chassis, you must enable clustering for all 3 module slots, or for container instances, a container instance in each slot, even if you do not have a module installed. If you do not configure all 3 modules, the cluster will not come up.

Before you begin

- Download the application image you want to use for the logical device from Cisco.com, and then upload that image to the Firepower 4100/9300 chassis.
- For container instances, if you do not want to use the default profile, add a resource profile according to Add a Resource Profile for Container Instances, on page 337.
- For container instances, before you can install a container instance for the first time, you must reinitialize the security module/engine so that the disk has the correct formatting. Choose **Security Modules** or **Security Engine**, and click the Reinitialize icon (). An existing logical device will be deleted and then reinstalled as a new device, losing any local application configuration. If you are replacing a native instance with container instances, you will need to delete the native instance in any case. You cannot automatically migrate a native instance to a container instance.
- Gather the following information:
  - Management interface ID, IP addresses, and network mask
  - Gateway IP address
  - Firewall Management Center IP address and/or NAT ID of your choosing
  - DNS server IP address
  - Firewall Threat Defense hostname and domain name

Procedure

Step 1 Configure interfaces.

a) Add at least one Data type interface or EtherChannel (also known as a port-channel) before you deploy the cluster. See Add an EtherChannel (Port Channel), on page 333 or Configure a Physical Interface, on page 332.

For clustering on multiple chassis, all data interfaces must be Spanned EtherChannels with at least one member interface. Add the same EtherChannels on each chassis. Combine the member interfaces from all cluster units into a single EtherChannel on the switch. See Clustering Guidelines and Limitations, on page 707 for more information about EtherChannels.

For multi-instance clustering, you cannot use FXOS-defined VLAN subinterfaces or data-sharing interfaces in the cluster. Only application-defined subinterfaces are supported. See FXOS Interfaces vs. Application Interfaces, on page 300 for more information.

b) Add a Management type interface or EtherChannel. See Add an EtherChannel (Port Channel), on page 333 or Configure a Physical Interface, on page 332.

The management interface is required. Note that this management interface is not the same as the chassis management interface that is used only for chassis management (in FXOS, you might see the chassis management interface displayed as MGMT, management0, or other similar names).

For clustering on multiple chassis, add the same Management interface on each chassis.

For multi-instance clustering, you can share the same management interface across multiple clusters on the same chassis, or with standalone instances.

c) For clustering on multiple chassis, add a member interface to the cluster control link EtherChannel (by default, port-channel 48). See Add an EtherChannel (Port Channel), on page 333.

Do not add a member interface for a cluster isolated to security modules within one Firepower 9300 chassis. If you add a member, the chassis assumes this cluster will be using multiple chassis, and will only allow you to use Spanned EtherChannels, for example.

On the **Interfaces** tab, the port-channel 48 cluster type interface shows the **Operation State** as **failed** if it does not include any member interfaces. For a cluster isolated to security modules within one Firepower 9300 chassis, this EtherChannel does not require any member interfaces, and you can ignore this Operational State.

Add the same member interfaces on each chassis. The cluster control link is a device-local EtherChannel on each chassis. Use separate EtherChannels on the switch per device. See Clustering Guidelines and Limitations, on page 707 for more information about EtherChannels.

For multi-instance clustering, you can create additional Cluster type EtherChannels. Unlike the Management interface, the cluster control link is *not* sharable across multiple devices, so you will need a Cluster interface for each cluster. However, we recommend using VLAN subinterfaces instead of multiple EtherChannels; see the next step to add a VLAN subinterface to the Cluster interface.

d) For multi-instance clustering, add VLAN subinterfaces to the cluster EtherChannel so you have a subinterface for each cluster. See Add a VLAN Subinterface for Container Instances, on page 335.

If you add subinterfaces to a Cluster interface, you cannot use that interface for a native cluster.

e) (Optional) Add an eventing interface. See Add an EtherChannel (Port Channel), on page 333 or Configure a Physical Interface, on page 332.

This interface is a secondary management interface for the Firewall Threat Defense devices. To use this interface, you must configure its IP address and other parameters at the Firewall Threat Defense CLI. For example, you can separate management traffic from events (such as web events). See the **configure network** commands in the Firewall Threat Defense command reference.

For clustering on multiple chassis, add the same eventing interface on each chassis.

Step 2 Choose **Logical Devices**.

Step 3 Click **Add** > **Cluster**, and set the following parameters:

Figure 301: Native Cluster

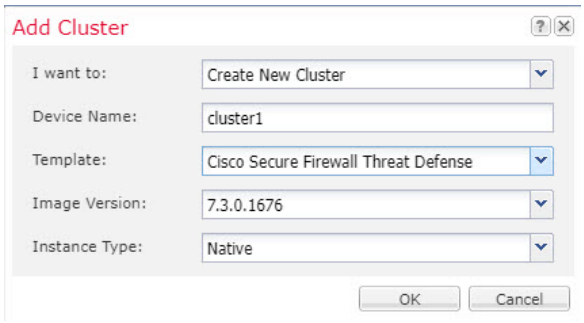

Figure 302: Multi-Instance Cluster

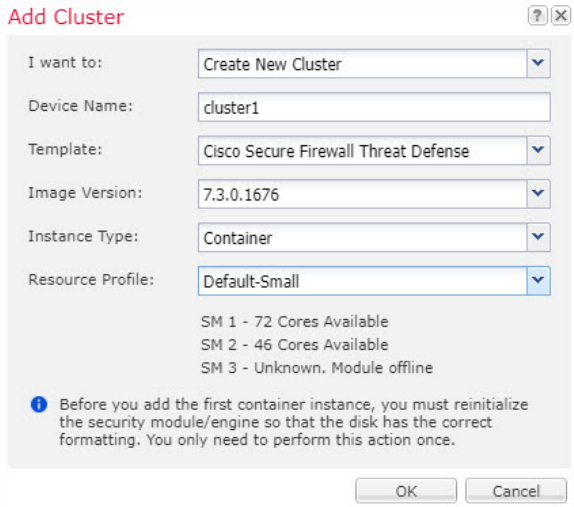

a) Choose **I want to:** > **Create New Cluster**
b) Provide a **Device Name**.

   This name is used internally by the chassis supervisor to configure management settings and to assign interfaces; it is not the device name used in the application configuration.

c) For the **Template**, choose **Cisco Firepower Threat Defense**.
d) Choose the **Image Version**.
e) For the **Instance Type**, choose either **Native** or **Container**.

   A native instance uses all of the resources (CPU, RAM, and disk space) of the security module/engine, so you can only install one native instance. A container instance uses a subset of resources of the security module/engine, so you can install multiple container instances.

f) (Container Instance only) For the **Resource Type**, choose one of the resource profiles from the drop-down list.

   For the Firepower 9300, this profile will be applied to each instance on each security module. You can set different profiles per security module later in this procedure; for example, if you are using different security module types, and you want to use more CPUs on a lower-end model. We recommend choosing the correct profile before you create the cluster. If you need to create a new profile, cancel out of the cluster creation, and add one using Add a Resource Profile for Container Instances, on page 337.

   Note
   If you assign a different profile to instances in an established cluster, which allows mismatched profiles, then apply the new profile on the data nodes first; after they reboot and come back up, you can apply the new profile to the control node.

g) Click **OK**.

   You see the Provisioning - *device name* window.

Step 4 Choose the interfaces you want to assign to this cluster.

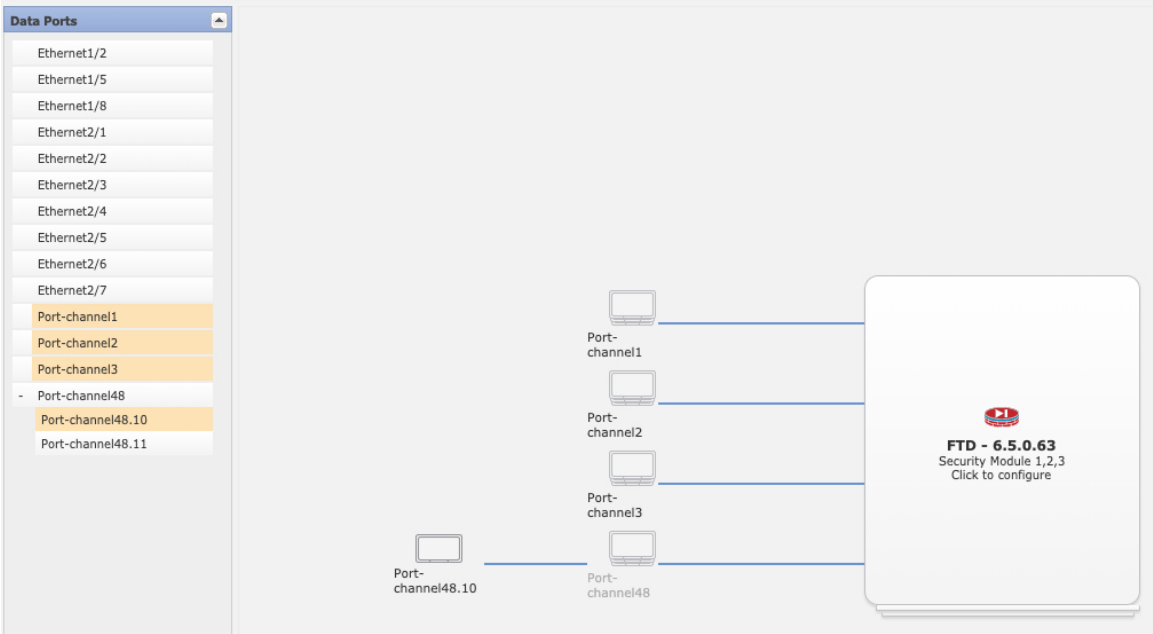

For native mode clustering: All valid interfaces are assigned by default. If you defined multiple Cluster type interfaces, deselect all but one.

For multi-instance clustering: Choose each data interface you want to assign to the cluster, and also choose the Cluster type port-channel or port-channel subinterface.

Step 5 Click the device icon in the center of the screen.

A dialog box appears where you can configure initial bootstrap settings. These settings are meant for initial deployment only, or for disaster recovery. For normal operation, you can later change most values in the application CLI configuration.

Step 6 On the **Cluster Information** page, complete the following.

Figure 303: Native Cluster

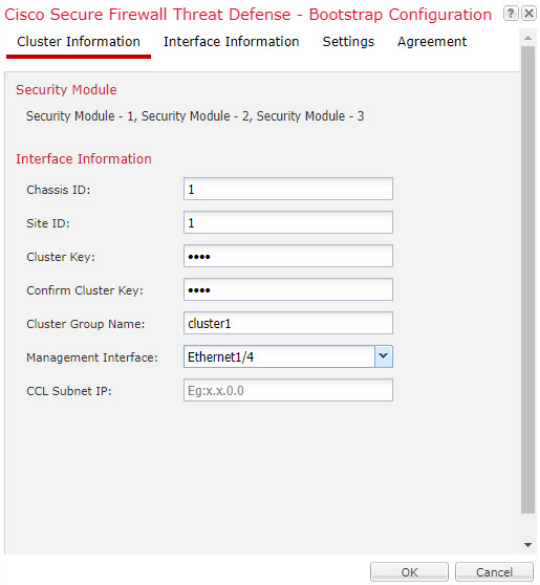

Figure 304: Multi-Instance Cluster

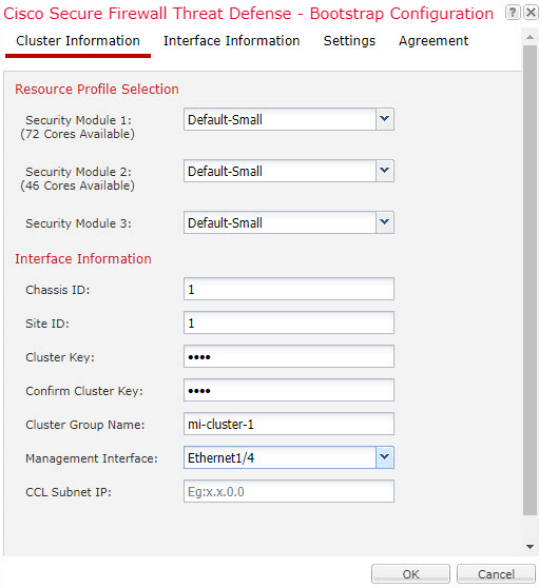

a) (Container Instance for the Firepower 9300 only) In the **Security Module (SM) and Resource Profile Selection** area, you can set a different resource profile per module; for example, if you are using different security module types, and you want to use more CPUs on a lower-end model.

b) For clustering on multiple chassis, in the **Chassis ID** field, enter a chassis ID. Each chassis in the cluster must use a unique ID.

   This field only appears if you added a member interface to cluster control link Port-Channel 48.

c) For inter-site clustering, in the **Site ID** field, enter the site ID for this chassis between 1 and 8. FlexConfig feature. Additional inter-site cluster customizations to enhance redundancy and stability, such as director

localization, site redundancy, and cluster flow mobility, are only configurable using the Firewall Management Center FlexConfig feature.

d) In the **Cluster Key** field, configure an authentication key for control traffic on the cluster control link.

The shared secret is an ASCII string from 1 to 63 characters. The shared secret is used to generate the key. This option does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear.

e) Set the **Cluster Group Name**, which is the cluster group name in the logical device configuration.

The name must be an ASCII string from 1 to 38 characters.

Important
From 2.4.1, spaces in cluster group name will be considered as special characters and may result in error while deploying the logical devices. To avoid this issue, you must rename the cluster group name without a space.

f) Choose the **Management Interface**.

This interface is used to manage the logical device. This interface is separate from the chassis management port.

If you assign a Hardware Bypass-capable interface as the Management interface, you see a warning message to make sure your assignment is intentional.

g) (Optional) Set the **CCL Subnet IP** as ***a.b*.0.0**.

By default, the cluster control link uses the 127.2.0.0/16 network. However, some networking deployments do not allow 127.2.0.0/16 traffic to pass. In this case, specify any /16 network address on a unique network for the cluster, except for loopback (127.0.0.0/8), multicast (224.0.0.0/4), and internal (169.254.0.0/16) addresses. If you set the value to 0.0.0.0, then the default network is used.

The chassis auto-generates the cluster control link interface IP address for each unit based on the chassis ID and slot ID: *a.b.chassis_id.slot_id*.

Step 7 On the **Settings** page, complete the following.

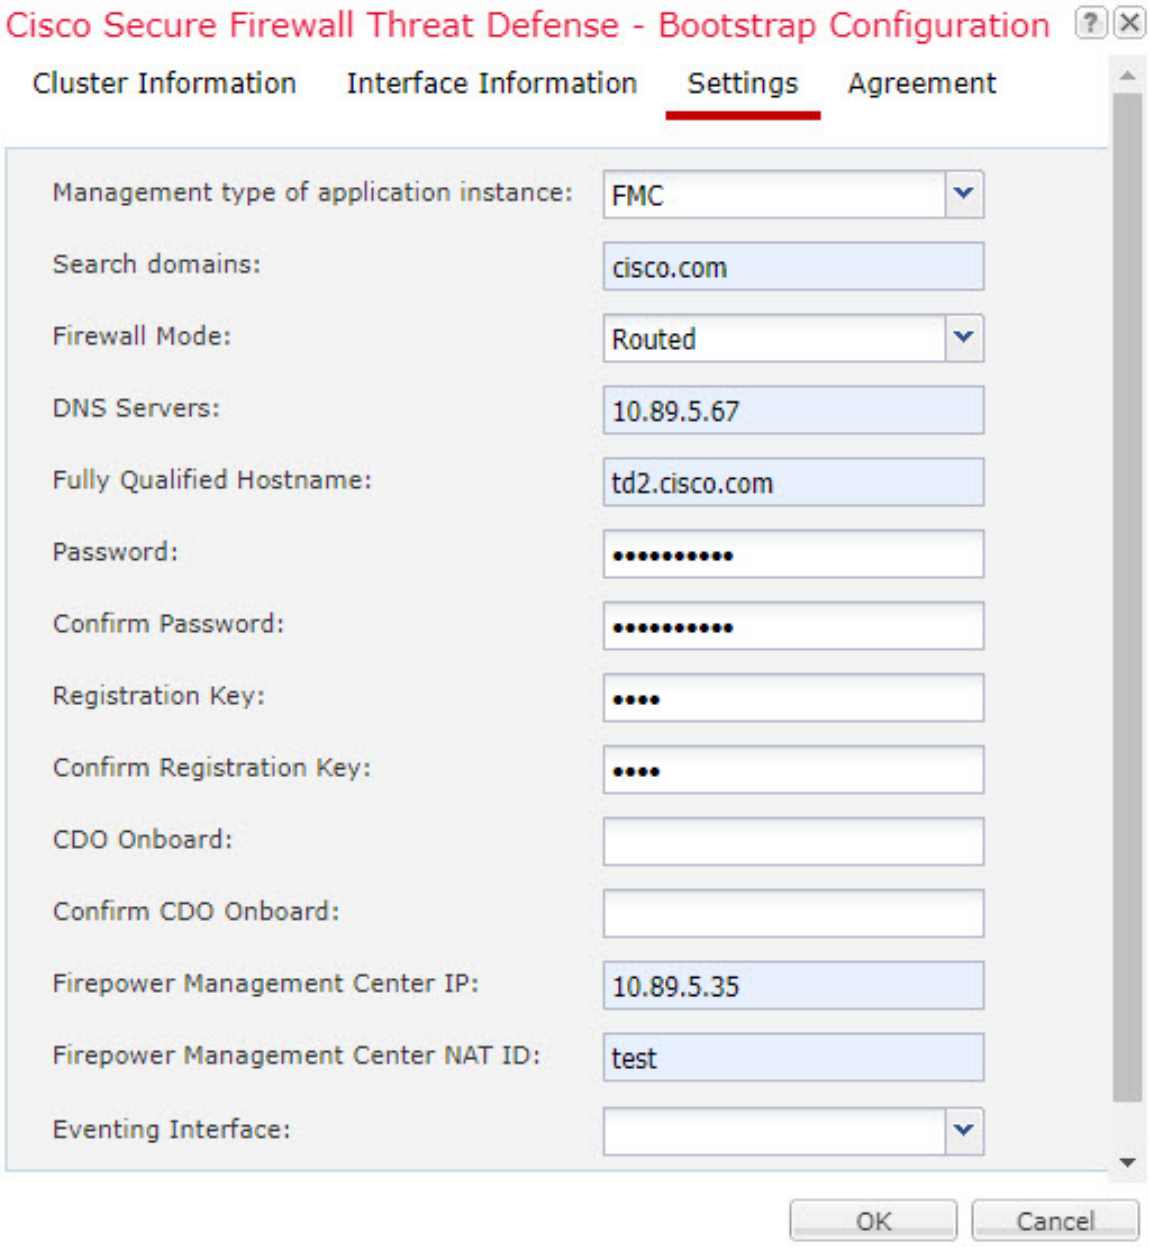

a) In the **Registration Key** field, enter the key to be shared between the Firewall Management Center and the cluster members during registration.

   You can choose any text string for this key between 1 and 37 characters; you will enter the same key on the Firewall Management Center when you add the Firewall Threat Defense.

b) Enter a **Password** for the Firewall Threat Defense admin user for CLI access.

c) In the **Firepower Management Center IP** field, enter the IP address of the managing Firewall Management Center. If you do not know the Firewall Management Center IP address, leave this field blank and enter a passphrase in the **Firepower Management Center NAT ID** field.

d) (Optional) For a container instance, **Permit Expert mode from FTD SSH sessions**: **Yes** or **No**. Expert Mode provides Firewall Threat Defense shell access for advanced troubleshooting.

If you choose **Yes** for this option, then users who access the container instance directly from an SSH sesssion can enter Expert Mode. If you choose **No**, then only users who access the container instance from the FXOS CLI can enter Expert Mode. We recommend choosing **No** to increase isolation between instances.

Use Expert Mode only if a documented procedure tells you it is required, or if the Cisco Technical Assistance Center asks you to use it. To enter this mode, use the **expert** command in the Firewall Threat Defense CLI.

e) (Optional) In the **Search Domains** field, enter a comma-separated list of search domains for the management network.

f) (Optional) From the **Firewall Mode** drop-down list, choose **Transparent** or **Routed**.

In routed mode, the Firewall Threat Defense is considered to be a router hop in the network. Each interface that you want to route between is on a different subnet. A transparent firewall, on the other hand, is a Layer 2 firewall that acts like a "bump in the wire," or a "stealth firewall," and is not seen as a router hop to connected devices.

The firewall mode is only set at initial deployment. If you re-apply the bootstrap settings, this setting is not used.

g) (Optional) In the **DNS Servers** field, enter a comma-separated list of DNS servers.

The Firewall Threat Defense uses DNS if you specify a hostname for the Firewall Management Center, for example.

h) (Optional) In the **Firepower Management Center NAT ID** field, enter a passphrase that you will also enter on the Firewall Management Center when you add the cluster as a new device.

Normally, you need both IP addresses (along with a registration key) for both routing purposes and for authentication: the Firewall Management Center specifies the device IP address, and the device specifies the Firewall Management Center IP address. However, if you only know one of the IP addresses, which is the minimum requirement for routing purposes, then you must also specify a unique NAT ID on both sides of the connection to establish trust for the initial communication and to look up the correct registration key. You can specify any text string as the NAT ID, from 1 to 37 characters. The Firewall Management Center and device use the registration key and NAT ID (instead of IP addresses) to authenticate and authorize for initial registration.

i) (Optional) In the **Fully Qualified Hostname** field, enter a fully qualified name for the Firewall Threat Defense device.

Valid characters are the letters from a to z, the digits from 0 to 9, the dot (.), and the hyphen (-); maximum number of characters is 253.

j) (Optional) From the **Eventing Interface** drop-down list, choose the interface on which events should be sent. If not specified, the management interface will be used.

To specify a separate interface to use for events, you must configure an interface as a *firepower-eventing* interface. If you assign a Hardware Bypass-capable interface as the Eventing interface, you see a warning message to make sure your assignment is intentional.

Step 8 On the **Interface Information** page, configure a management IP address for each security module in the cluster. Select the type of address from the **Address Type** drop-down list and then complete the following for each security module.

Note
You must set the IP address for all 3 module slots in a chassis, even if you do not have a module installed. If you do not configure all 3 modules, the cluster will not come up.

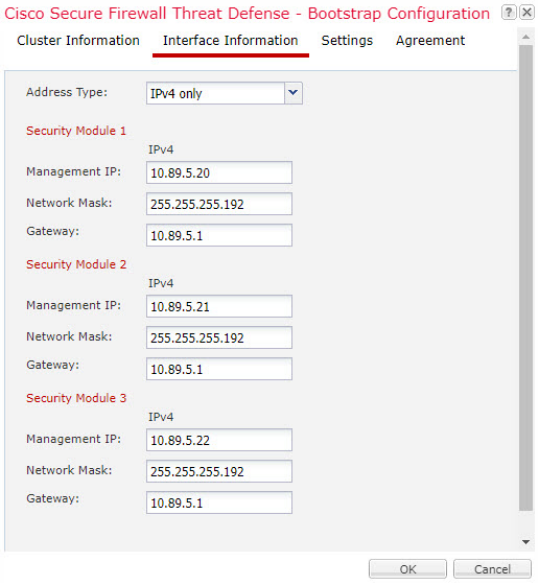

a) In the **Management IP** field, configure an IP address.

   Specify a unique IP address on the same network for each module.

b) Enter a **Network Mask** or **Prefix Length**.

c) Enter a **Network Gateway** address.

Step 9 On the **Agreement** tab, read and accept the end user license agreement (EULA).

Step 10 Click **OK** to close the configuration dialog box.

Step 11 Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for the status of the new logical device. When the logical device shows its **Status** as **online**, you can add the remaining cluster chassis, or for a cluster isolated to security modules within one Firepower 9300 chassis, start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.

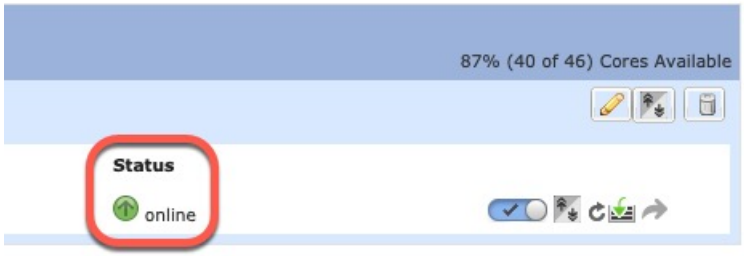

Step 12 For clustering on multiple chassis, add the next chassis to the cluster:

a) On the first chassis of the Firewall Chassis Manager, click the **Show Configuration icon** at the top right; copy the displayed cluster configuration.

b) Connect to the Firewall Chassis Manager on the next chassis, and add a logical device according to this procedure.

c) Choose **I want to: > Join an Existing Cluster**.

d) Click **OK**.

e) In the **Copy Cluster Details** box, paste in the cluster configuration from the first chassis, and click **OK**.

f) Click the device icon in the center of the screen. The cluster information is mostly pre-filled, but you must change the following settings:

- **Chassis ID**—Enter a unique chassis ID.
- **Site ID**—For inter-site clustering, enter the site ID for this chassis between 1 and 8. Additional inter-site cluster customizations to enhance redundancy and stability, such as director localization, site redundancy, and cluster flow mobility, are only configurable using the Firewall Management Center FlexConfig feature.
- **Cluster Key**—(Not prefilled) Enter the same cluster key.
- **Management IP**—Change the management address for each module to be a unique IP address on the same network as the other cluster members.

Click **OK**.

g) Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for each cluster member for the status of the new logical device. When the logical device for each cluster member shows its **Status** as **online**, you can start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.

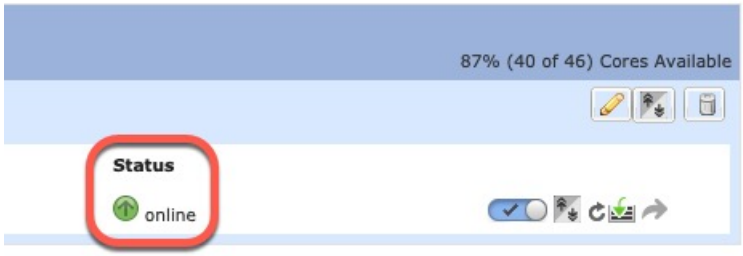

Step 13 Add the control unit to the Firewall Management Center using the management IP address.

All cluster units must be in a successfully-formed cluster on FXOS prior to adding them to Firewall Management Center.

The Firewall Management Center then automatically detects the data units.

Add More Cluster Nodes

Add or replace the Firewall Threat Defense cluster node in an existing cluster. When you add a new cluster node in FXOS, the Firewall Management Center adds the node automatically.

Note The FXOS steps in this procedure only apply to adding a new *chassis*; if you are adding a new module to a Firepower 9300 where clustering is already enabled, the module will be added automatically.

Before you begin

- In the case of a replacement, you must delete the old cluster node from the Firewall Management Center. When you replace it with a new node, it is considered to be a new device on the Firewall Management Center.
- The interface configuration must be the same on the new chassis. You can export and import FXOS chassis configuration to make this process easier.

Procedure

Step 1 If you previously upgraded the Firewall Threat Defense image using the Firewall Management Center, perform the following steps *on each chassis in the cluster*.

When you upgraded from the Firewall Management Center, the startup version in the FXOS configuration was not updated, and the standalone package was not installed on the chassis. Both of these items need to be set manually so the new node can join the cluster using the correct image version.

Note
If you only applied a patch release, you can skip this step. Cisco does not provide standalone packages for patches.

a) Install the running Firewall Threat Defense image on the chassis using the **System** > **Updates** page.
b) Click **Logical Devices** and click the Set Version icon ( ). For a Firepower 9300 with multiple modules, set the version for each module.

The **Startup Version** shows the original package you deployed with. The **Current Version** shows the version you upgraded to.

c) In the **New Version** drop-down menu, choose the version that you uploaded. This version should match the **Current Version** displayed, and will set the startup version to match the new version.
d) On the new chassis, make sure the new image package is installed.

Step 2 On an existing cluster chassis Firewall Chassis Manager, click **Logical Devices**.

Step 3 Click the **Show Configuration icon** at the top right; copy the displayed cluster configuration.

Step 4 Connect to the Firewall Chassis Manager on the new chassis, and click **Add** > **Cluster**.

Step 5 For the **Device Name**, provide a name for the logical device.

Step 6 Click **OK**.

Step 7 In the **Copy Cluster Details** box, paste in the cluster configuration from the first chassis, and click **OK**.

Step 8 Click the device icon in the center of the screen. The cluster information is partly pre-filled, but you must fill in the following settings:

Figure 305: Cluster Information

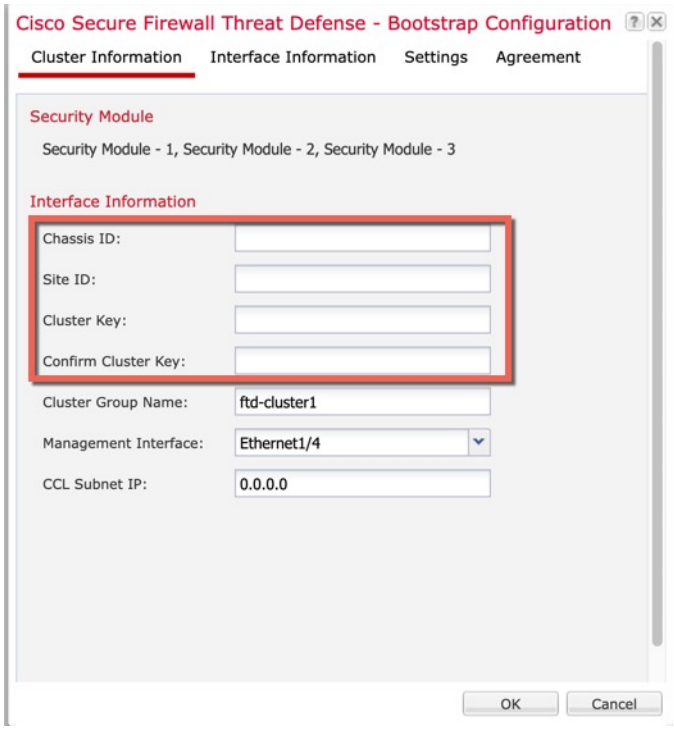

Figure 306: Interface Information

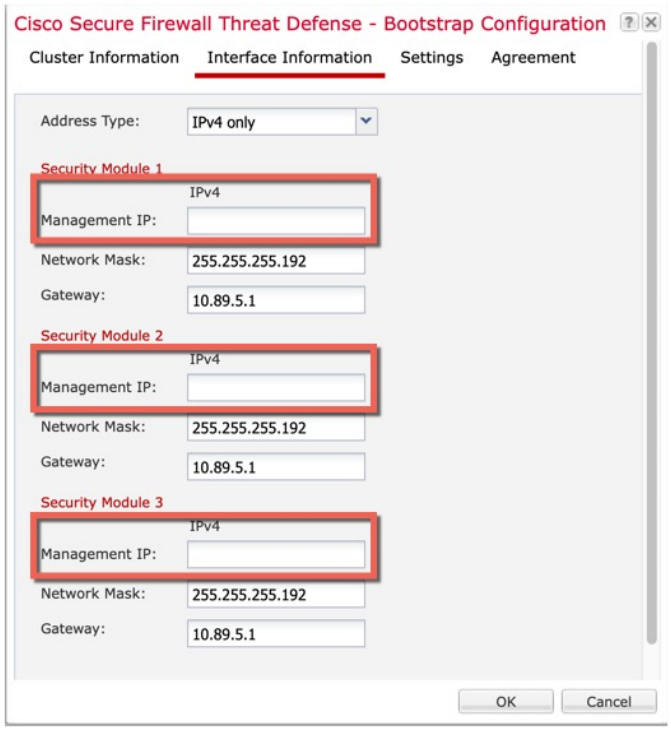

Figure 307: Settings

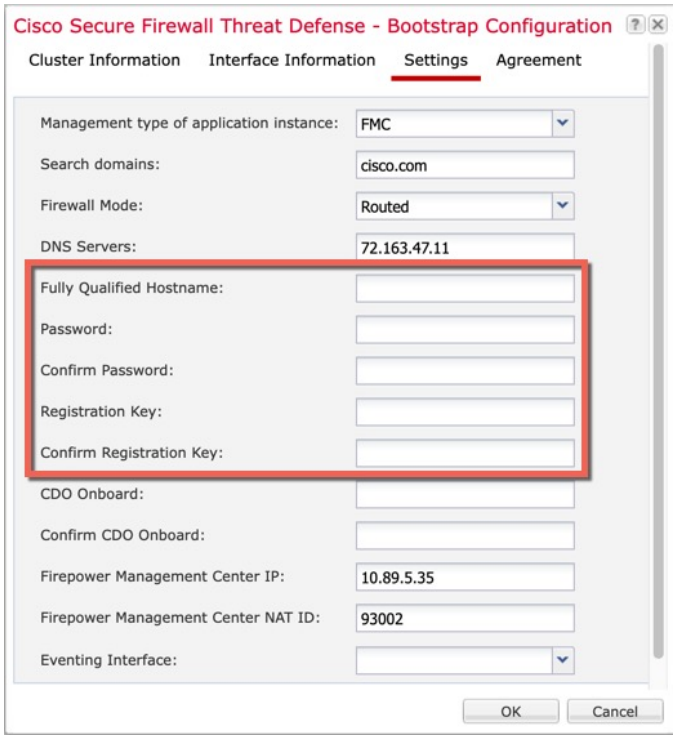

- **Chassis ID**—Enter a *unique* chassis ID.

- **Site ID**—For inter-site clustering, enter the site ID for this chassis between 1 and 8. This feature is only configurable using the Firewall Management Center FlexConfig feature.
- **Cluster Key**—Enter the *same* cluster key.
- **Management IP**—Change the management address for each module to be a *unique* IP address on the same network as the other cluster members.
- **Fully Qualified Hostname**—Enter the *same* hostname.
- **Password**—Enter the *same* password.
- **Registration Key**—Enter the *same* registration key.

Click **OK**.

Step 9 Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for each cluster member for the status of the new logical device. When the logical device for each cluster member shows its **Status** as **online**, you can start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.

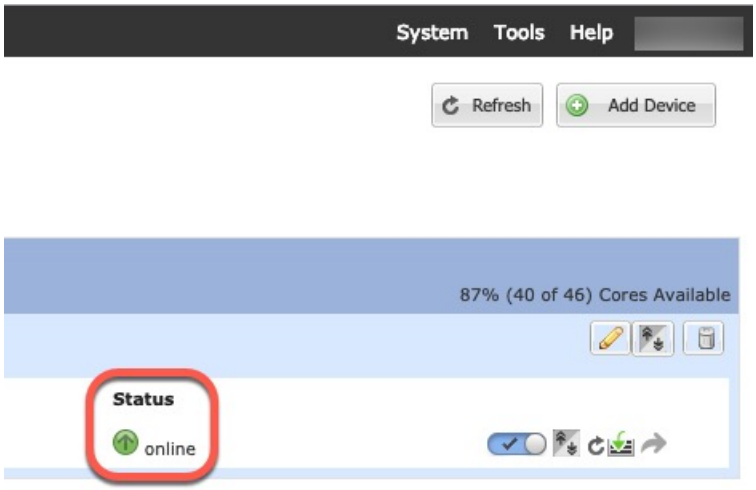

Firewall Management Center: Add a Cluster

Add one of the cluster units as a new device to the Secure Firewall Management Center; the Firewall Management Center auto-detects all other cluster members.

Before you begin

- All cluster units must be in a successfully-formed cluster on FXOS prior to adding the cluster to the Firewall Management Center. You should also check which unit is the control unit. Refer to the Firewall Chassis Manager **Logical Devices** screen or use the Firewall Threat Defense **show cluster info** command.

Procedure

Step 1 In the Firewall Management Center, choose **Devices** > **Device Management**, and then choose **Add** > **Add Device** to add the control unit using the unit's management IP address you assigned when you deployed the cluster.

Figure 308: Add Device

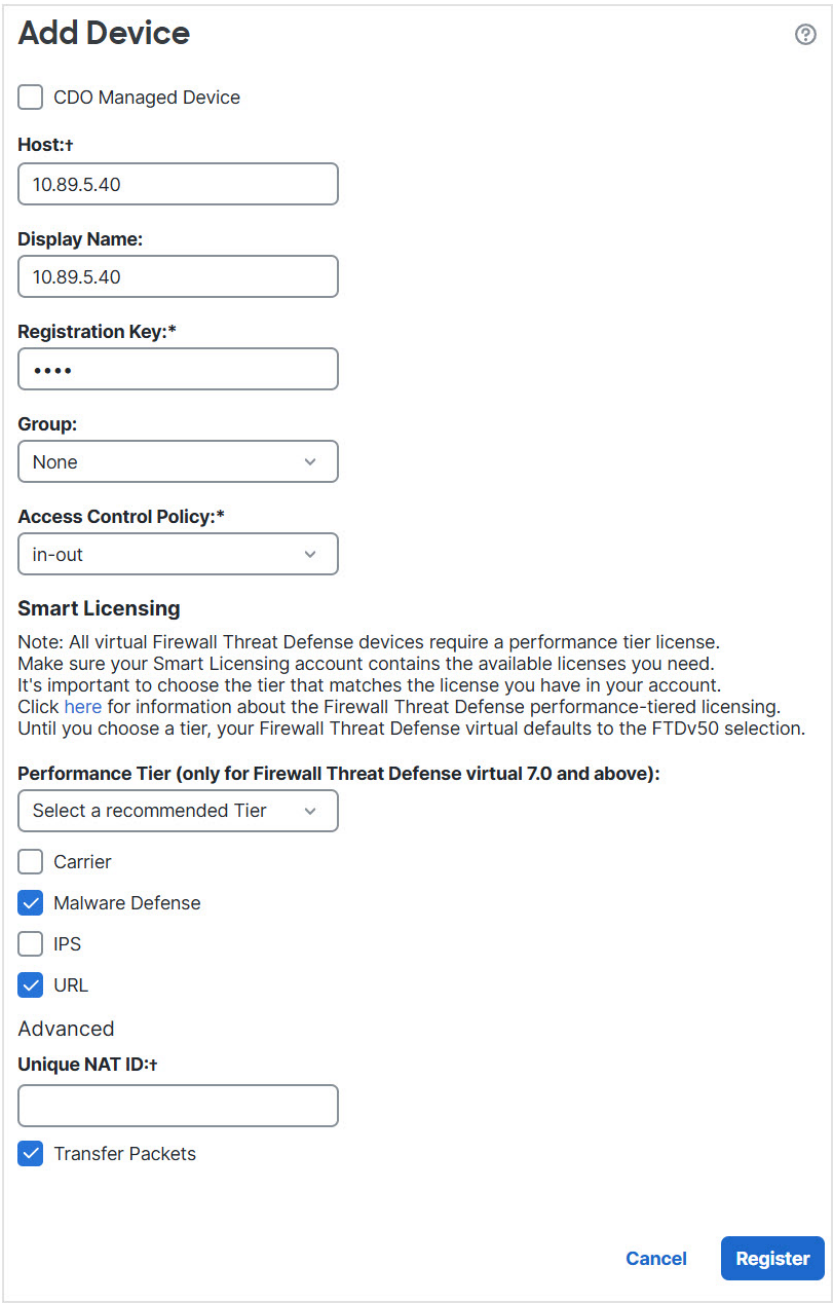

a) In the **Host** field, enter the IP address or hostname of the control unit.

We recommend adding the control unit for the best performance, but you can add any unit of the cluster.

If you used a NAT ID during device setup, you may not need to enter this field. For more information, see NAT Environments, on page 7.

b) In the **Display Name** field, enter a name for the control unit as you want it to display in the Firewall Management Center.

This display name is not for the cluster; it is only for the control unit you are adding. You can later change the name of other cluster members and the cluster display name.

c) In the **Registration Key** field, enter the same registration key that you used when you deployed the cluster in FXOS. The registration key is a one-time-use shared secret.

d) (Optional) Add the device to a device **Group**.

e) Choose an initial **Access Control Policy** to deploy to the device upon registration, or create a new policy.

If you create a new policy, you create a basic policy only. You can later customize the policy as needed.

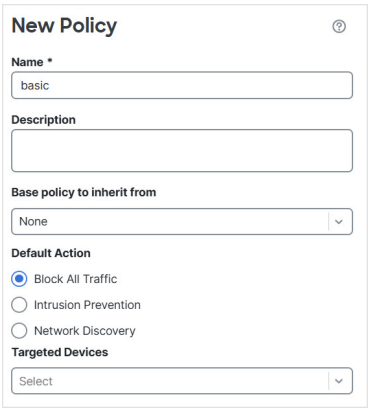

f) Choose licenses to apply to the device.

g) If you used a NAT ID during device setup, expand the **Advanced** section and enter the same NAT ID in the **Unique NAT ID** field.

h) Check the **Transfer Packets** check box to allow the device to transfer packets to the Firewall Management Center.

This option is enabled by default. When events like IPS or Snort are triggered with this option enabled, the device sends event metadata information and packet data to the Firewall Management Center for inspection. If you disable it, only event information will be sent to the Firewall Management Center but packet data is not sent.

i) Click **Register**.

The Firewall Management Center identifies and registers the control unit, and then registers all data units. If the control unit does not successfully register, then the cluster is not added. A registration failure can occur if the cluster was not up on the chassis, or because of other connectivity issues. In this case, we recommend that you try re-adding the cluster unit.

The cluster name shows on the **Devices** > **Device Management** page; expand the cluster to see the cluster units.

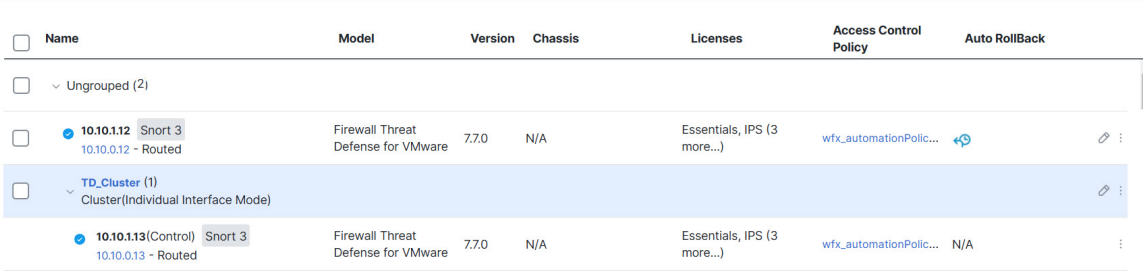

A unit that is currently registering shows the loading icon.

You can monitor cluster unit registration by clicking the **Notifications** icon and choosing **Tasks**. The Firewall Management Center updates the Cluster Registration task as each unit registers. If any units fail to register, see Reconcile Cluster Members, on page 744.

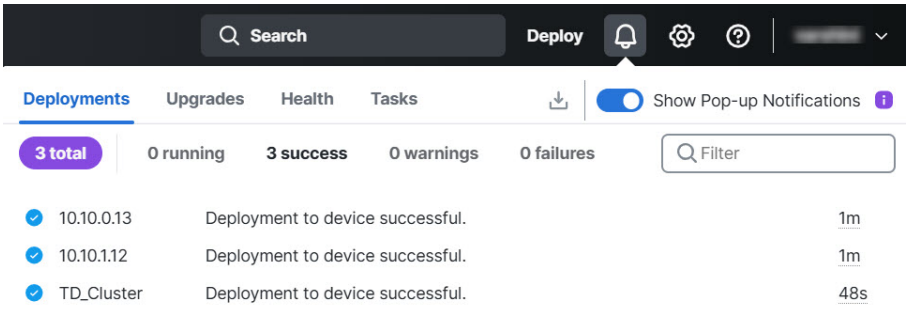

Step 2 Configure device-specific settings by clicking the **Edit** (✎) for the cluster.

Most configuration can be applied to the cluster as a whole, and not member units in the cluster. For example, you can change the display name per unit, but you can only configure interfaces for the whole cluster.

Step 3 On the **Devices** > **Device Management** > **Cluster** screen, you see **General**, **License**, **System**, and **Health** settings.

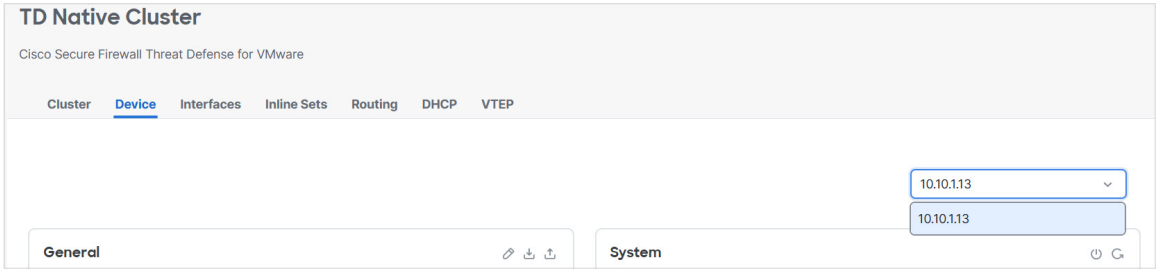

See the following cluster-specific items:

- **General** > **Name**—Change the cluster display name by clicking the **Edit** (✎).

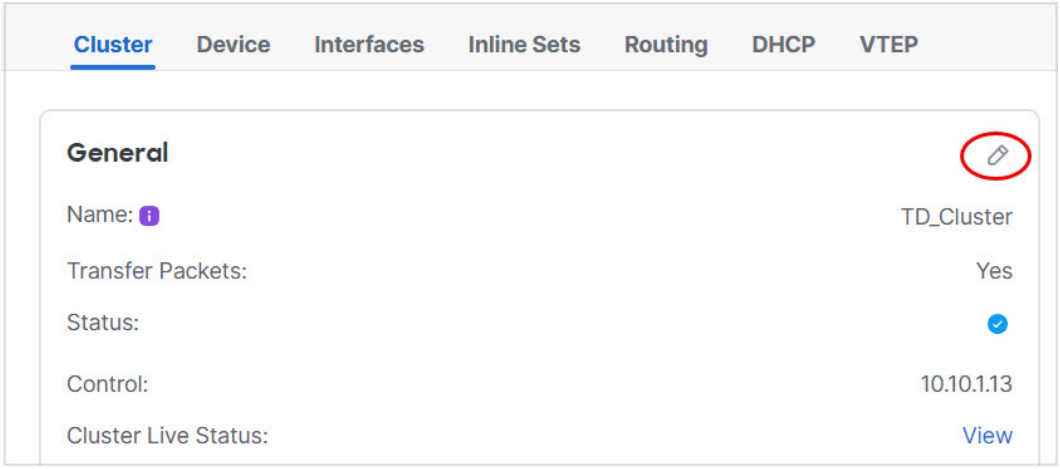

Then set the **Name** field.

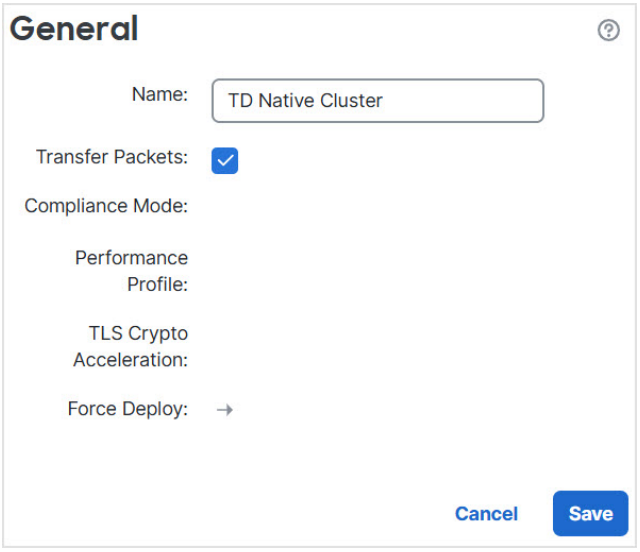

- **General > View cluster status**—Click the **View cluster status** link to open the **Cluster Status** dialog box.

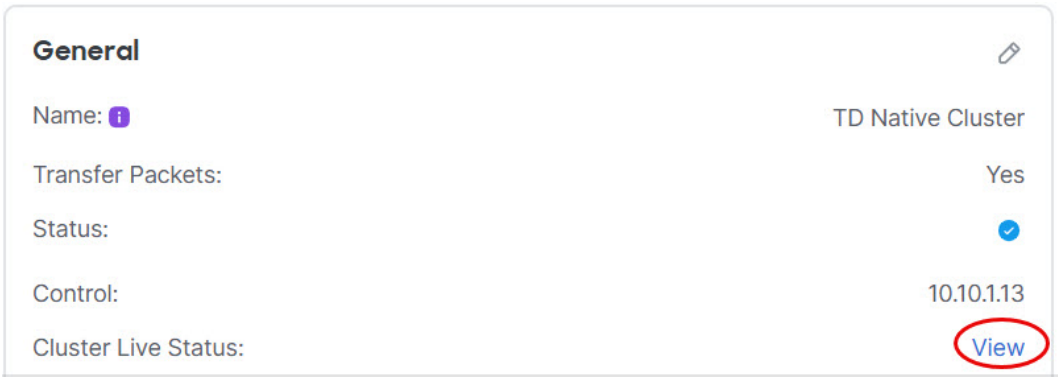

The **Cluster Status** dialog box also lets you retry data unit registration by clicking **Reconcile**.

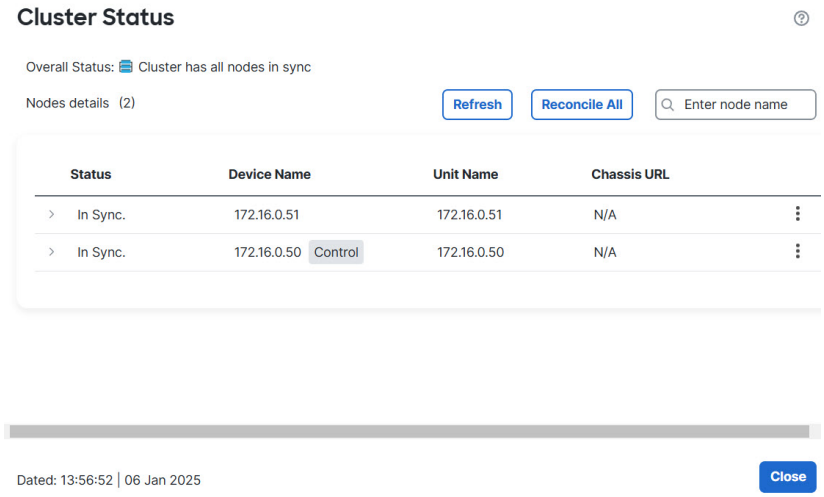

- **General > Troubleshoot**—You can generate and download troubleshooting logs, and you can view cluster CLIs. See Firewall Management Center: Troubleshooting the Cluster, on page 749.

Figure 309: Troubleshoot

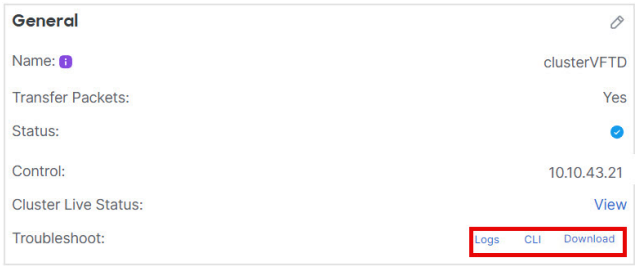

- **License**—Click **Edit** (✎) to set license entitlements.

Step 4 On the **Devices** > **Device Management** > **Devices**, you can choose each member in the cluster from the top right drop-down menu and configure the following settings.

- **General > Name**—Change the cluster member display name by clicking the **Edit** (✎).

Then set the **Name** field.

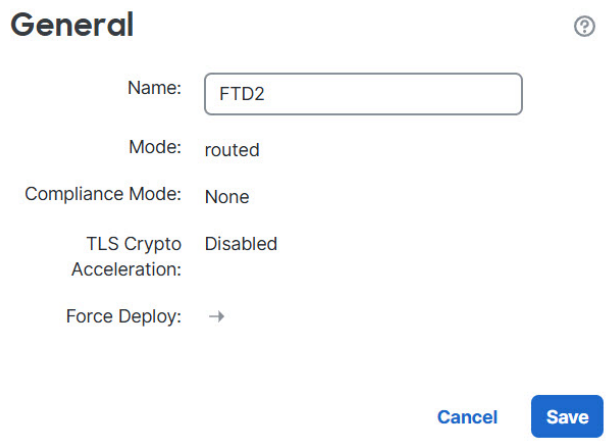

- **Management > Host**—If you change the management IP address in the device configuration, you must match the new address in the Firewall Management Center so that it can reach the device on the network; edit the **Host** address in the **Management** area.

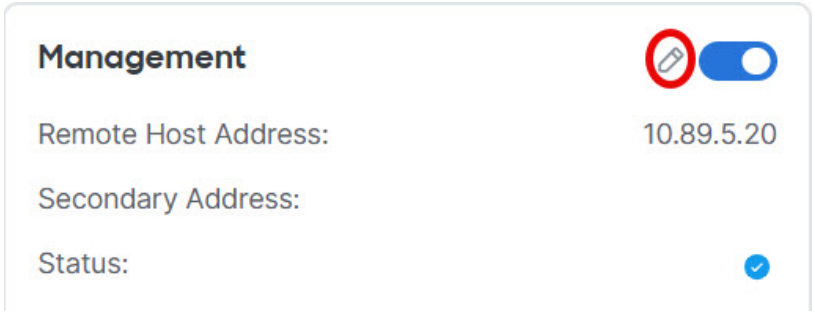

Firewall Management Center: Configure Cluster, Data Interfaces

This procedure configures basic parameters for each data interface that you assigned to the cluster when you deployed it in FXOS. For clustering on multiple chassis, data interfaces are always Spanned EtherChannel interfaces. For the cluster control link interface for a cluster isolated to security modules within one Firepower 9300 chassis, you must increase the MTU from the default.

Note When using Spanned EtherChannels for clustering on multiple chassis, the port-channel interface will not come up until clustering is fully enabled. This requirement prevents traffic from being forwarded to a unit that is not an active unit in the cluster.

Procedure

Step 1 Choose **Devices** > **Device Management**, and click **Edit** (✎) next to the cluster.

Step 2 Click **Interfaces**.

Step 3 Configure the cluster control link.

For clustering on multiple chassis, set the cluster control link MTU to be at least 100 bytes higher than the highest MTU of the data interfaces. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead. We suggest setting the MTU to the maximum of 9184; the minimum value is 1400 bytes. For example, because the maximum MTU is 9184, then the highest data interface MTU can be 9084, while the cluster control link can be set to 9184.

When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the ping fails, a notification is generated so you can fix the MTU mismatch on connecting switches and try again.

For native clusters: The cluster control link interface is Port-Channel48 by default. If you don't know which interface is the cluster control link, check the FXOS configuration for chassis for the Cluster-type interface assigned to the cluster.

a) Click **Edit** (✎) for the cluster control link interface.
b) On the **General** page, in the **MTU** field, enter a value between 1400 and 9184but not between 2561 and 8362. Due to block pool handling, this MTU size is not optimal for system operation. We suggest using the maximum, 9184.
c) Click **OK**.

Step 4 Configure data interfaces.

a) (Optional) For regular firewall interfaces, configure VLAN subinterfaces on the data interface. The rest of this procedure applies to the subinterfaces. See Add a Subinterface, on page 836.
b) Click **Edit** (✎) for the data interface.
c) Configure the name and other parameters. For regular firewall interfaces, see Configure Routed Mode Interfaces, on page 856 or, for transparent mode, Configure Bridge Group Interfaces, on page 861. For IPS-only interfaces, see Inline Sets and Passive Interfaces, on page 893..

Note
If the cluster control link interface MTU is not at least 100 bytes higher than the data interface MTU, you will see an error that you must reduce the MTU of the data interface. See, Step 3, on page 732 to increase the cluster control link MTU, after which you can continue configuring the data interfaces.

d) For clustering on multiple chassis, set a unique. manual global MAC address for the EtherChannel. Click **Advanced**, and in the **Active Mac Address** field, enter a MAC address in H.H.H format, where H is a 16-bit hexadecimal digit.

For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE. The MAC address must not have the multicast bit set, that is, the second hexadecimal digit from the left cannot be an odd number.

Do not set the **Standby Mac Address**; it is ignored.

You must configure a unique MAC address not currently in use on your network for a Spanned EtherChannel to avoid potential network connectivity problems. With a manually-configured MAC address, the MAC address stays with the current control unit. If you do not configure a MAC address, then if the control unit changes, the new control unit uses a new MAC address for the interface, which can cause a temporary network outage.

e) Click **OK**. Repeat the above steps for other data interfaces.

Step 5 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Firewall Management Center: Configure Cluster Health Monitor Settings

The **Cluster Health Monitor Settings** section of the **Cluster** page displays the settings described in the table below.

Figure 310: Cluster Health Monitor Settings

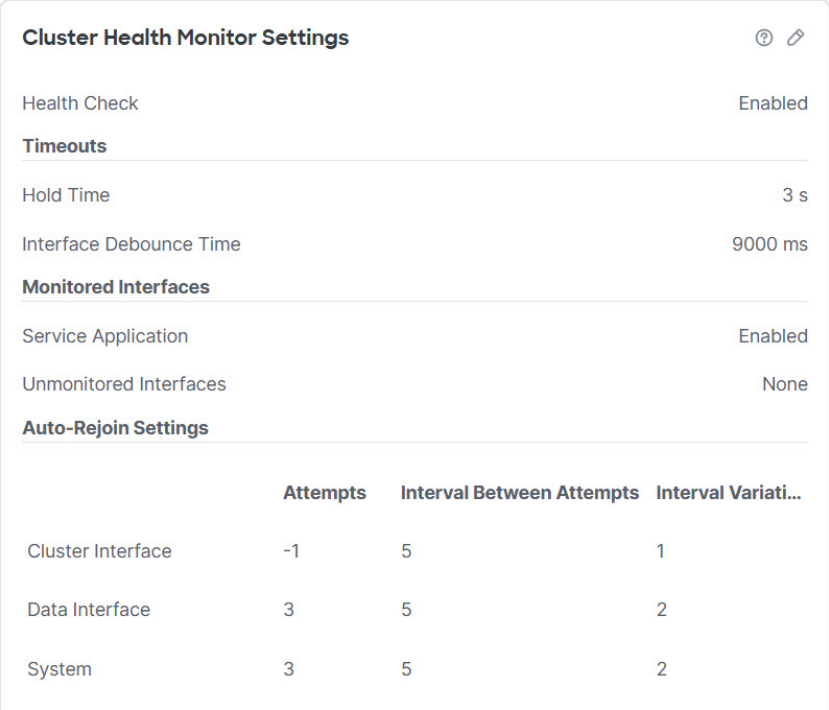

Table 46: Cluster Health Monitor Settings Section Table Fields

Field	Description
Timeouts	
Hold Time	Between .3 and 45 seconds; The default is 3 seconds. To determine node system health, the cluster nodes send heartbeat messages on the cluster control link to other nodes. If a node does not receive any heartbeat messages from a peer node within the hold time period, the peer node is considered unresponsive or dead.
Interface Debounce Time	Between 300 and 9000 ms. The default is 500 ms. The interface debounce time is the amount of time before the node considers an interface to be failed, and the node is removed from the cluster.

Field	Description
Monitored Interfaces	The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The amount of time before the node removes a member from the cluster depends on the type of interface and whether the node is an established node or is joining the cluster.
Service Application	Shows whether the Snort and disk-full processes are monitored.
Unmonitored Interfaces	Shows unmonitored interfaces.
Auto-Rejoin Settings	
Cluster Interface	Shows the auto-rejoin settings after a cluster control link failure.
Attempts	Between -1 and 65535. The default is -1 (unlimited). Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 1x the interval duration. Defines if the interval duration increases at each attempt.
Data Interfaces	Shows the auto-rejoin settings after a data interface failure.
Attempts	Between -1 and 65535. The default is 3. Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 2x the interval duration. Defines if the interval duration increases at each attempt.
System	Shows the auto-rejoin settings after internal errors. Internal failures include: application sync timeout; inconsistent application statuses; and so on.
Attempts	Between -1 and 65535. The default is 3. Sets the number of rejoin attempts.
Interval Between Attempts	Between 2 and 60. The default is 5 minutes. Defines the interval duration in minutes between rejoin attempts.
Interval Variation	Between 1 and 3. The default is 2x the interval duration. Defines if the interval duration increases at each attempt.

Note If you disable the system health check, fields that do not apply when the system health check is disabled will not show.

You can change these settings from this section.

You can monitor any port-channel ID, single physical interface ID, as well as the Snort and disk-full processes. Health monitoring is not performed on VLAN subinterfaces or virtual interfaces such as VNIs or BVIs. You cannot configure monitoring for the cluster control link; it is always monitored.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Next to the cluster you want to modify, click **Edit** ( ).

Step 3 Click **Cluster**.

Step 4 In the **Cluster Health Monitor Settings** section, click **Edit** ( ).

Step 5 Disable the system health check by clicking the **Health Check** slider .

Figure 311: Disable the System Health Check

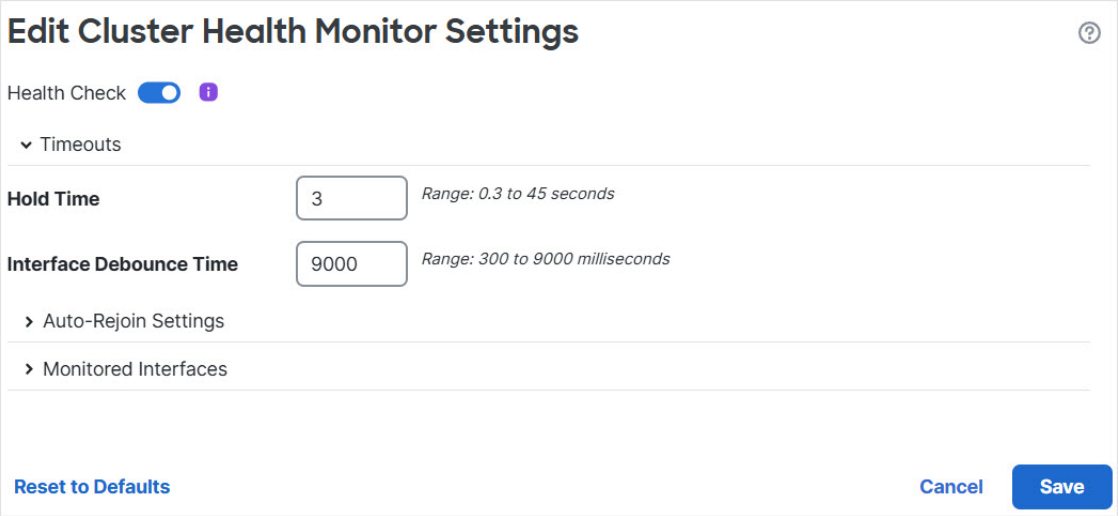

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 6 Configure the hold time and interface debounce time.

- **Hold Time**—Set the hold time to determine the amount of time between node heartbeat status messages, between .3 and 45 seconds; The default is 3 seconds.
- **Interface Debounce Time**—Set the debounce time between 300 and 9000 ms. The default is 500 ms. Lower values allow for faster detection of interface failures. Note that configuring a lower debounce time increases the chances of false-positives. When an interface status update occurs, the node waits the number of milliseconds specified before marking the interface as failed, and the node is removed from the cluster. In the case of an EtherChannel that transitions from a down state to an up state (for example, the switch reloaded, or the switch enabled an EtherChannel), a longer debounce time can prevent the interface from appearing to be failed on a cluster node just because another cluster node was faster at bundling the ports.

Step 7 Customize the auto-rejoin cluster settings after a health check failure.

Figure 312: Configure Auto-Rejoin Settings

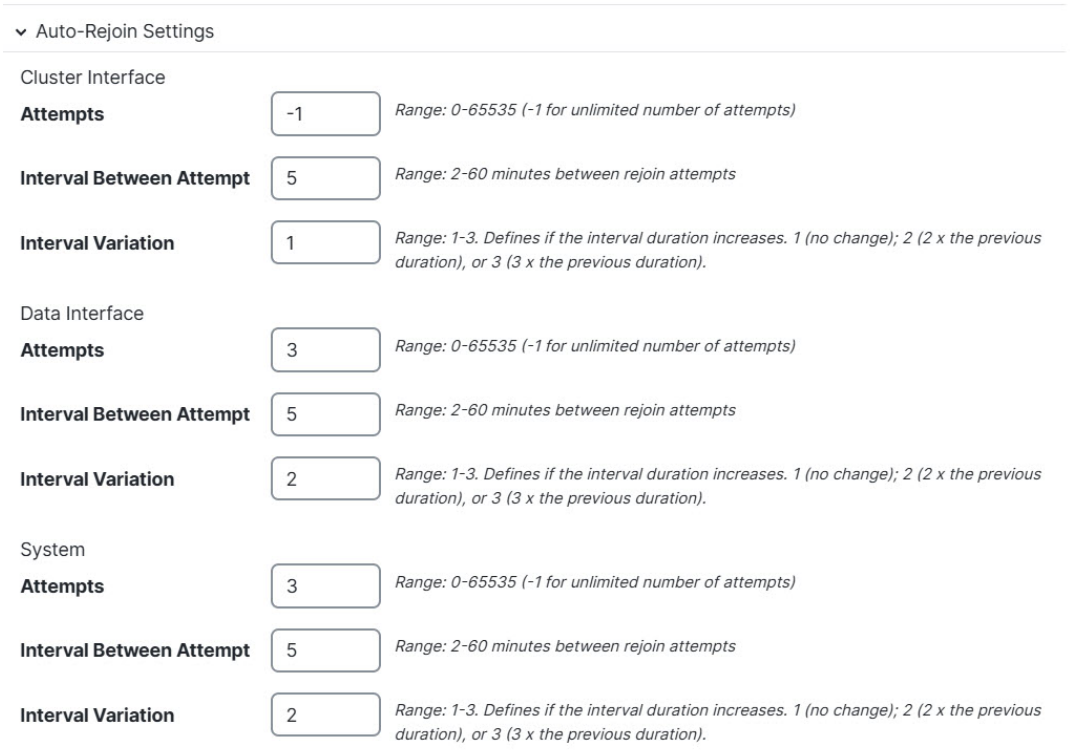

Set the following values for the **Cluster Interface**, **Data Interface**, and **System** (internal failures include: application sync timeout; inconsistent application statuses; and so on):

- **Attempts**—Sets the number of rejoin attempts, between -1 and 65535. **0** disables auto-rejoining. The default for the **Cluster Interface** is -1 (unlimited). The default for the **Data Interface** and **System** is 3.
- **Interval Between Attempts**—Defines the interval duration in minutes between rejoin attempts, between 2 and 60. The default value is 5 minutes. The maximum total time that the node attempts to rejoin the cluster is limited to 14400 minutes (10 days) from the time of last failure.
- **Interval Variation**—Defines if the interval duration increases. Set the value between 1 and 3: **1** (no change); **2** (2 x the previous duration), or **3** (3 x the previous duration). For example, if you set the interval duration to 5 minutes, and set the variation to 2, then the first attempt is after 5 minutes; the 2nd attempt is 10 minutes (2 x 5); the 3rd attempt 20 minutes (2 x 10), and so on. The default value is **1** for the **Cluster Interface** and **2** for the **Data Interface** and **System**.

Step 8 Configure monitored interfaces by moving interfaces in the **Monitored Interfaces** or **Unmonitored Interfaces** window. You can also check or uncheck **Enable Service Application Monitoring** to enable or disable monitoring of the Snort and disk-full processes.

Figure 313: Configure Monitored Interfaces

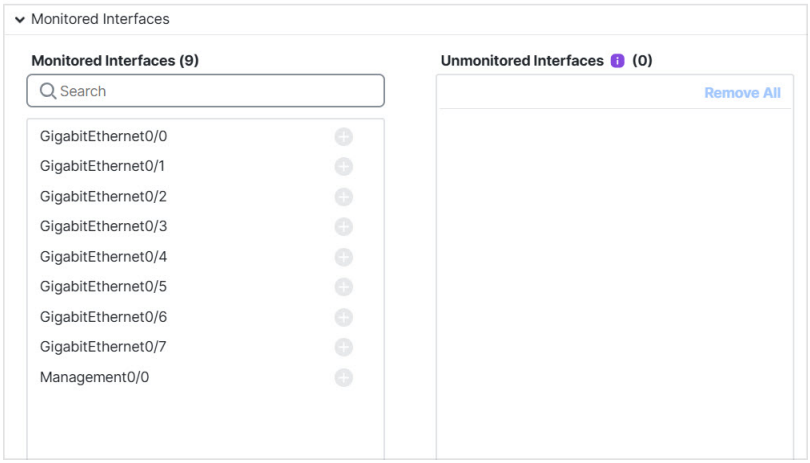

The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The amount of time before the node removes a member from the cluster depends on the type of interface and whether the node is an established node or is joining the cluster. Health check is enabled by default for all interfaces and for the Snort and disk-full processes.

You might want to disable health monitoring of non-essential interfaces.

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 9 Click **Save**.

Step 10 Deploy configuration changes; see Deploy Configuration Changes, on page 259.

FXOS: Remove a Cluster Node

The following sections describe how to remove nodes temporarily or permanently from the cluster.

Temporary Removal

A cluster node will be automatically removed from the cluster due to a hardware or network failure, for example. This removal is temporary until the conditions are rectified, and it can rejoin the cluster. You can also manually disable clustering.

To check whether a device is currently in the cluster, check the cluster status on the Firewall Chassis Manager **Logical Devices** page:

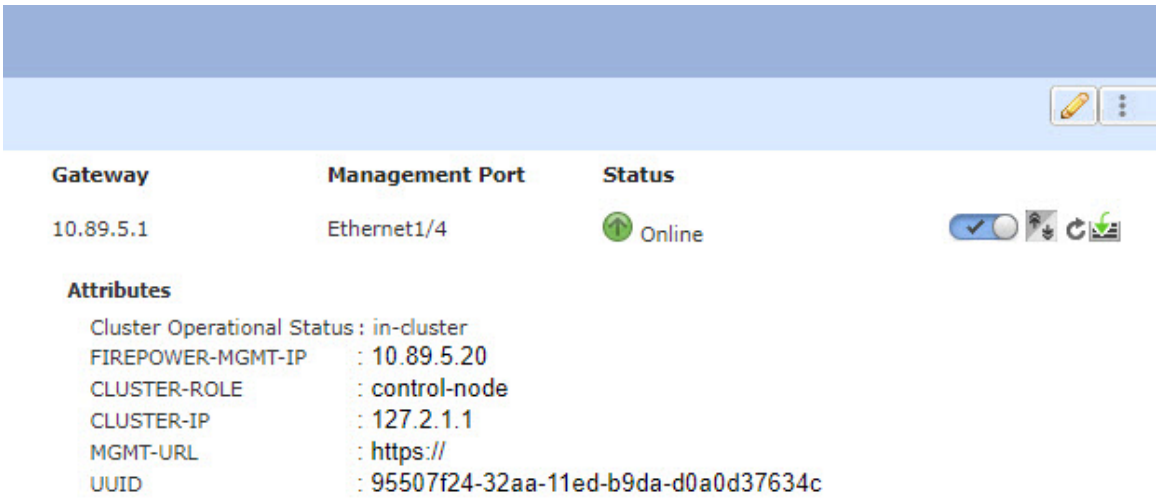

For Firewall Threat Defense using the Firewall Management Center, you should leave the device in the Firewall Management Center device list so that it can resume full functionality after you reenable clustering.

- Disable clustering in the application—You can disable clustering using the application CLI. Enter the **cluster remove unit** *name* command to remove any node other than the one you are logged into. The bootstrap configuration remains intact, as well as the last configuration synced from the control node, so you can later re-add the node without losing your configuration. If you enter this command on a data node to remove the control node, a new control node is elected.

  When a device becomes inactive, all data interfaces are shut down; only the Management interface can send and receive traffic. To resume traffic flow, re-enable clustering. The Management interface remains up using the IP address the node received from the bootstrap configuration. However if you reload, and the node is still inactive in the cluster , the Management interface is disabled.

  To reenable clustering, on the Firewall Threat Defense enter **cluster enable**.

- Disable the application instance—In the Firewall Chassis Manager on the **Logical Devices** page, click the **Slider enabled** ( ). You can later reenable it using the **Slider disabled** ( ).

- Shut down the security module/engine—In the Firewall Chassis Manager on the **Security Module/Engine** page, click the **Power Off icon**.

- Shut down the chassis—In the Firewall Chassis Manager on the **Overview** page, click the **Shut Down icon**.

Permanent Removal

You can permanently remove a cluster node using the following methods.

For Firewall Threat Defense using the Firewall Management Center, be sure to remove the node from the Firewall Management Center device list after you disable clustering on the chassis.

- Delete the logical device—In the Firewall Chassis Manager on the **Logical Devices** page, click the **Delete** ( ). You can then deploy a standalone logical device, a new cluster, or even add a new logical device to the same cluster.

- Remove the chassis or security module from service—If you remove a device from service, you can add replacement hardware as a new node of the cluster.

Firewall Management Center: Manage Cluster Members

After you deploy the cluster, you can change the configuration and manage cluster members.

Add a New Cluster Member

When you add a new cluster member in FXOS, the Secure Firewall Management Center adds the member automatically.

Before you begin

- Make sure the interface configuration is the same on the replacement unit as for the other chassis.

Procedure

Step 1 Add the new unit to the cluster in FXOS. See the FXOS configuration guide.

Wait for the new unit to be added to the cluster. Refer to the Firewall Chassis Manager **Logical Devices** screen or use the Firewall Threat Defense **show cluster info** command to view cluster status.

Step 2 The new cluster member is added automatically. To monitor the registration of the replacement unit, view the following:

- **Cluster Status** dialog box (which is available from the **Devices** > **Device Management More** (⋮) icon or from the **Devices** > **Device Management** > **Cluster** tab > **General** area > **Cluster Live Status** link)—A unit that is joining the cluster on the chassis shows "Joining cluster..." After it joins the cluster, the Firewall Management Center attempts to register it, and the status changes to "Available for Registration". After it completes registration, the status changes to "In Sync." If the registration fails, the unit will stay at "Available for Registration". In this case, force a re-registration by clicking **Reconcile**.
- System status > **Tasks** —The Firewall Management Center shows all registration events and failures.
- **Devices** > **Device Management**—When you expand the cluster on the devices listing page, you can see when a unit is registering when it has the loading icon to the left.

Replace a Cluster Member

You can replace a cluster member in an existing cluster. The Firewall Management Center auto-detects the replacement unit. However, you must manually delete the old cluster member in the Firewall Management Center. This procedure also applies to a unit that was reinitialized; in this case, although the hardware remains the same, it appears to be a new member.

Before you begin

- Make sure the interface configuration is the same on the replacement unit as for other chassis.

Procedure

Step 1 For a new chassis, if possible, backup and restore the configuration from the old chassis in FXOS.

If you are replacing a module in a Firepower 9300, you do not need to perform these steps.

If you do not have a backup FXOS configuration from the old chassis, first perform the steps in Add a New Cluster Member, on page 739.

For information about all of the below steps, see the FXOS configuration guide.

a) Use the configuration export feature to export an XML file containing logical device and platform configuration settings for your Firepower 4100/9300 chassis.
b) Import the configuration file to the replacement chassis.
c) Accept the license agreement.
d) If necessary, upgrade the logical device application instance version to match the rest of the cluster.

Step 2 In the Firewall Management Center for the old unit, choose **Devices** > **Device ManagementMore** (⋮) > **Delete** .

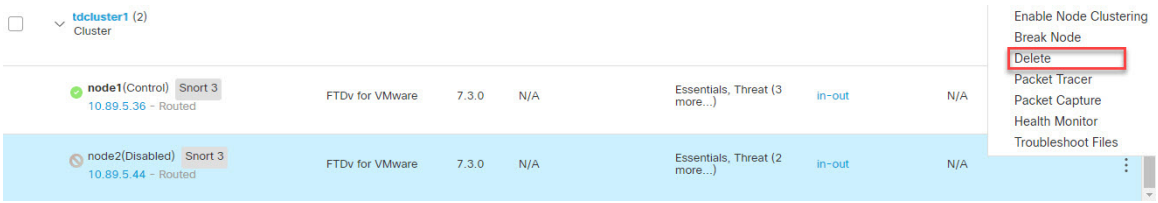

Step 3 Confirm that you want to delete the unit.

The unit is removed from the cluster and from the Firewall Management Center devices list.

Step 4 The new or reinitialized cluster member is added automatically. To monitor the registration of the replacement unit, view the following:

- **Cluster Status** dialog box (**Devices** > **Device Management More** (⋮) icon or **Devices** > **Device Management** > **Cluster** page**General** area **Cluster Live Status** link)—A unit that is joining the cluster on the chassis shows "Joining cluster..." After it joins the cluster, the Firewall Management Center attempts to register it, and the status changes to "Available for Registration". After it completes registration, the status changes to "In Sync." If the registration fails, the unit will stay at "Available for Registration". In this case, force a re-registration by clicking **Reconcile All**.
- **System** (⚙) > **Tasks**—The Firewall Management Center shows all registration events and failures.
- **Devices** > **Device Management**—When you expand the cluster on the devices listing page, you can see when a unit is registering when it has the loading icon to the left.

Deactivate a Member

You may want to deactivate a member in preparation for deleting the unit, or temporarily for maintenance. This procedure is meant to temporarily deactivate a member; the unit will still appear in the Firewall Management Center device list.

Note When a unit becomes inactive, all data interfaces are shut down; only the Management interface can send and receive traffic. To resume traffic flow, reenable clustering. The Management interface remains up using the IP address the unit received from the bootstrap configuration. However if you reload, and the unit is still inactive in the cluster, the management interface is disabled. You must use the console for any further configuration.

Procedure

Step 1 For the unit you want to deactivate, choose **Devices** > **Device Management More** (⋮) **Disable Clustering**.

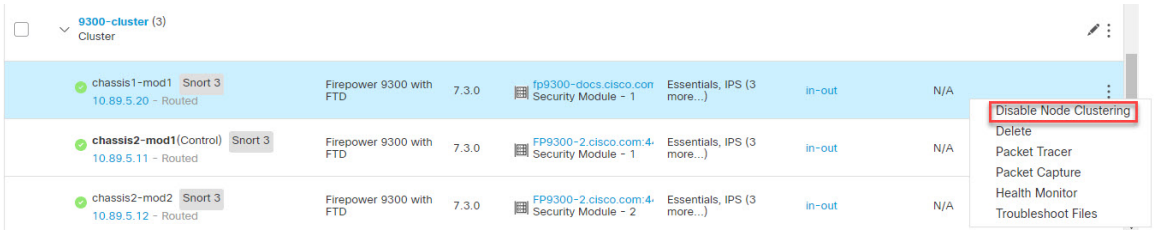

You can also deactivate a unit from the **Cluster Status** dialog box (**Devices** > **Device Management More** (⋮) **Cluster Live Status**).

Step 2 Confirm that you want to disable clustering on the unit.

The unit will show **(Disabled)** next to its name in the **Devices** > **Device Management** list.

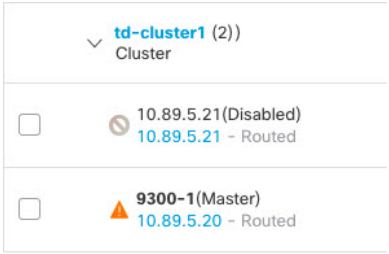

Step 3 To reenable clustering, see Rejoin the Cluster, on page 741.

Rejoin the Cluster

If a unit was removed from the cluster, for example for a failed interface or if you manually disabled clustering, you must manually rejoin the cluster. Make sure the failure is resolved before you try to rejoin the cluster. See Rejoining the Cluster, on page 760 for more information about why a unit can be removed from a cluster.

Procedure

Step 1 For the unit you want to reactivate, choose **Devices** > **Device Management More** (⋮) > **Enable Clustering**.

You can also reactivate a unit from the **Cluster Status** dialog box (**Devices** > **Device Management More** (⋮) > **Cluster Live Status**).

Step 2 Confirm that you want to enable clustering on the unit.

Unregister a Data Node

If you need to permanently remove a cluster node (for example, if you remove a module on the Firepower 9300, or remove a chassis), then you should unregister it from the Firewall Management Center.

Do not unregister the node if it is still a healthy part of the cluster, or if you only want to disable the node temporarily. To remove it permanently from the cluster in FXOS, see FXOS: Remove a Cluster Node, on page 737. If you unregister it from the Firewall Management Center, and it is still part of the cluster, it will continue to pass traffic, and could even become the control node—a control node that the Firewall Management Center can no longer manage.

Before you begin

To manually deactivate the node, see Deactivate a Member, on page 741. Before you unregister a node, the node must be inactive, either manually or because of a health failure.

Procedure

Step 1 Make sure the node is ready to be unregistered from the Firewall Management Center. On **Devices** > **Device Management**, make sure the node shows **(Disabled)**.

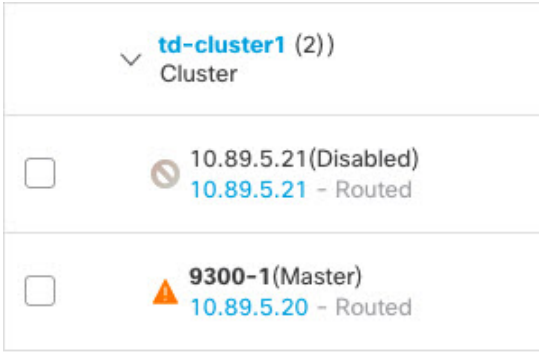

You can also view each node's status on the **Cluster Status** dialog box available from **More** (⋮). If the status is stale, click **Reconcile All** on the **Cluster Status** dialog box to force an update.

Step 2 In the Firewall Management Center for the data node you want to delete, choose **Devices** > **Device ManagementMore** (⋮) > **Unregister**.

Step 3 Confirm that you want to unregister the node.

The node is removed from the cluster and from the Firewall Management Center devices list.

Change the Control Unit

⚠

Caution The best method to change the control unit is to disable clustering on the control unit, wait for a new control election, and then re-enable clustering. If you must specify the *exact* unit you want to become the control unit, use the procedure in this section. Note that for centralized features, if you force a control unit change, then all connections are dropped, and you have to re-establish the connections on the new control unit.

To change the control unit, perform the following steps.

Procedure

Step 1 Open the **Cluster Status** dialog box by choosing **Devices** > **Device Management More** (⋮) **Cluster Live Status**.

You can also access the **Cluster Status** dialog box from **Devices** > **Device Management** > **Cluster** page **General** area **Cluster Live Status** link.

Step 2 For the unit you want to become the control unit, choose **More** (⋮) **Change Role to Control**.

Step 3 You are prompted to confirm the role change. Check the checkbox, and click **OK**.

Reconcile Cluster Members

If a cluster member fails to register, you can reconcile the cluster membership from the chassis to the Secure Firewall Management Center. For example, a data unit might fail to register if the Firewall Management Center is occupied with certain processes, or if there is a network issue.

Procedure

Step 1 Choose **Devices** > **Device Management More** (⋮) for the cluster, and then choose **Cluster Live Status** to open the **Cluster Status** dialog box.

You can also open the **Cluster Status** dialog box from the **Devices** > **Device Management** page **General** area **Cluster Live Status** link.

Step 2 Click **Reconcile All**.

For more information about the cluster status, see Firewall Management Center: Monitoring the Cluster, on page 744.

Firewall Management Center: Monitoring the Cluster

You can monitor the cluster in Secure Firewall Management Center and at the Firewall Threat Defense CLI.

- **Cluster Status** dialog box, which is available from the **Devices** > **Device Management More** (⋮) icon or from the **Devices** > **Device Management** > **Cluster** page**General** area **>** **Cluster Live Status** link.

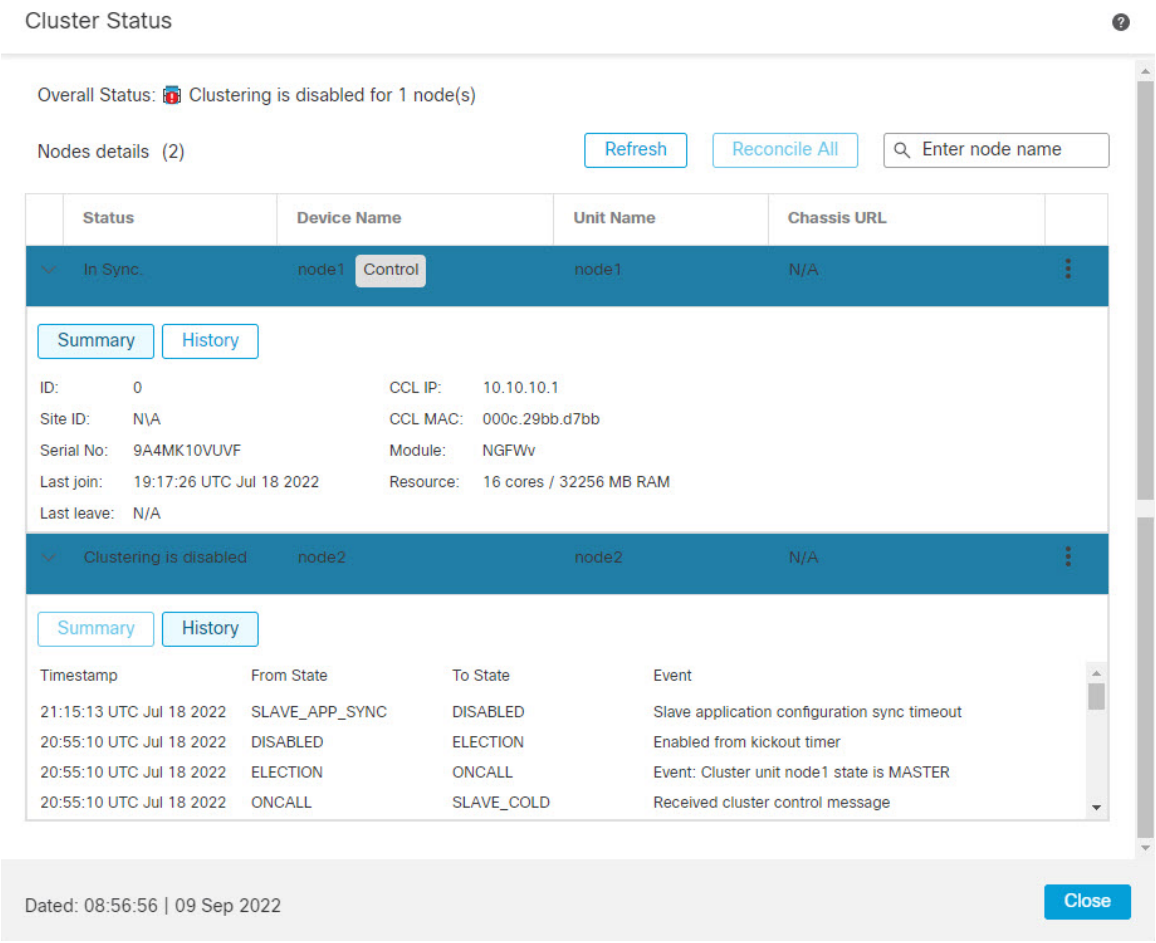

The Control unit has a graphic indicator identifying its role.

Cluster member **Status** includes the following states:

- In Sync.—The unit is registered with the Firewall Management Center.
- Pending Registration—The unit is part of the cluster, but has not yet registered with the Firewall Management Center. If a unit fails to register, you can retry registration by clicking **Reconcile All**.
- Clustering is disabled—The unit is registered with the Firewall Management Center, but is an inactive member of the cluster. The clustering configuration remains intact if you intend to later re-enable it, or you can delete the unit from the cluster.
- Joining cluster...—The unit is joining the cluster on the chassis, but has not completed joining. After it joins, it will register with the Firewall Management Center.

For each unit, you can view the **Summary** or the **History**.

For each unit from the **More** (⋮) menu , you can perform the following status changes:

- **Disable Clustering**
- **Enable Clustering**
- **Change Role to Control**

- **System** (⚙)**Tasks** page.

  The **Tasks** page shows updates of the Cluster Registration task as each unit registers.

- **Devices** > **Device Management** > *cluster_name*.

  When you expand the cluster on the devices listing page, you can see all member units, including the control unit shown with its role next to the IP address. For units that are still registering, you can see the loading icon.

- **show cluster** {**access-list** [*acl_name*] | **conn** [**count**] | **cpu** [**usage**] | **history** | **interface-mode** | **memory** | **resource usage** | **service-policy** | **traffic** | **xlate count**}

  To view aggregated data for the entire cluster or other information, use the **show cluster** command.

- **show cluster info** [**auto-join** | **clients** | **conn-distribution** | **flow-mobility counters** | **goid** [*options*] | **health** | **incompatible-config** | **loadbalance** | **old-members** | **packet-distribution** | **trace** [*options*] | **transport** { **asp** | **cp**}]

  To view cluster information, use the **show cluster info** command.

Cluster Health Monitor Dashboard

Cluster Health Monitor

When a Firewall Threat Defense is the control node of a cluster, the Firewall Management Center collects various metrics periodically from the device metric data collector. The cluster health monitor is comprised of the following components:

- Overview dashboard—Displays information about the cluster topology, cluster statistics, and metric charts:
  - The topology section displays a cluster's live status, the health of individual threat defense, threat defense node type (control node or data node), and the status of the device. The status of the device could be *Disabled* (when the device leaves the cluster), *Added out of box* (in a public cloud cluster, the additional nodes that do not belong to the Firewall Management Center), or *Normal* (ideal state of the node).
  - The cluster statistics section displays current metrics of the cluster with respect to the CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.

Note The CPU and memory metrics display the individual average of the data plane and snort usage.

  - The metric charts, namely, CPU Usage, Memory Usage, Throughput, and Connections, diagrammatically display the statistics of the cluster over the specified time period.

- Load Distribution dashboard—Displays load distribution across the cluster nodes in two widgets:
  - The Distribution widget displays the average packet and connection distribution over the time range across the cluster nodes. This data depicts how the load is being distributed by the nodes. Using this widget, you can easily identify any abnormalities in the load distribution and rectify it.

- The Node Statistics widget displays the node level metrics in table format. It displays metric data on CPU usage, memory usage, input rate, output rate, active connections, and NAT translations across the cluster nodes. This table view enables you to correlate data and easily identify any discrepancies.

- Member Performance dashboard—Displays current metrics of the cluster nodes. You can use the selector to filter the nodes and view the details of a specific node. The metric data include CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.
- CCL dashboard—Displays, graphically, the cluster control link data namely, the input, and output rate.
- Troubleshooting and Links — Provides convenient links to frequently used troubleshooting topics and procedures.
- Time range—An adjustable time window to constrain the information that appears in the various cluster metrics dashboards and widgets.
- Custom Dashboard—Displays data on both cluster-wide metrics and node-level metrics. However, node selection only applies for the threat defense metrics and not for the entire cluster to which the node belongs.

Viewing Cluster Health

You must be an Admin, Maintenance, or Security Analyst user to perform this procedure.

The cluster health monitor provides a detailed view of the health status of a cluster and its nodes. This cluster health monitor provides health status and trends of the cluster in an array of dashboards.

Before you begin

- Ensure you have created a cluster from one or more devices in the Firewall Management Center.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Use the Monitoring navigation pane to access node-specific health monitors.

Step 2 In the device list, click **Expand**(>) and **Collapse** (∨) to expand and collapse the list of managed cluster devices.

Step 3 To view the cluster health statistics, click on the cluster name. The cluster monitor reports health and performance metrics in several predefined dashboards by default. The metrics dashboards include:

- Overview — Highlights key metrics from the other predefined dashboards, including its nodes, CPU, memory, input and output rates, connection statistics, and NAT translation information.
- Load Distribution — Traffic and packet distribution across the cluster nodes.
- Member Performance — Node-level statistics on CPU usage, memory usage, input throughput, output throughput, active connection, and NAT translation.
- CCL — Interface status and aggregate traffic statistics.

You can navigate through the various metrics dashboards by clicking on the labels. For a comprehensive list of the supported cluster metrics, see Cisco Secure Firewall Threat Defense Health Metrics.

Step 4 You can configure the time range from the drop-down in the upper-right corner. The time range can reflect a period as short as the last hour (the default) or as long as two weeks. Select **Custom** from the drop-down to configure a custom start and end date.

Click the refresh icon to set auto refresh to 5 minutes or to toggle off auto refresh.

Step 5 Click on deployment icon for a deployment overlay on the trend graph, with respect to the selected time range.

The deployment icon indicates the number of deployments during the selected time-range. A vertical band indicates the deployment start and end time. For multiple deployments, multiple bands/lines appear. Click on the icon on top of the dotted line to view the deployment details.

Step 6 (For node-specific health monitor) View the **Health Alerts** for the node in the alert notification at the top of page, directly to the right of the device name.

Hover your pointer over the **Health Alerts** to view the health summary of the node. The popup window shows a truncated summary of the top five health alerts. Click on the popup to open a detailed view of the health alert summary.

Step 7 (For node-specific health monitor) The device monitor reports health and performance metrics in several predefined dashboards by default. The metrics dashboards include:

- Overview — Highlights key metrics from the other predefined dashboards, including CPU, memory, interfaces, connection statistics; plus disk usage and critical process information.
- CPU — CPU utilization, including the CPU usage by process and by physical cores.
- Memory — Device memory utilization, including data plane and Snort memory usage.
- Interfaces — Interface status and aggregate traffic statistics.
- Connections — Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
- Snort — Statistics that are related to the Snort process.
- ASP drops — Statistics related to the dropped packets against various reasons.

You can navigate through the various metrics dashboards by clicking on the labels. See Cisco Secure Firewall Threat Defense Health Metrics for a comprehensive list of the supported device metrics.

Step 8 Click the plus sign **Add New Dashboard**(+) in the upper right corner of the health monitor to create a custom dashboard by building your own variable set from the available metric groups.

For cluster-wide dashboard, choose Cluster metric group, and then choose the metric.

Cluster Metrics

The cluster health monitor tracks statistics that are related to a cluster and its nodes, and aggregate of load distribution, performance, and CCL traffic statistics.

Table 47: Cluster Metrics

Metric	Description	Format
CPU	Average of CPU metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Memory	Average of memory metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Data Throughput	Incoming and outgoing data traffic statistics for a cluster.	bytes
CCL Throughput	Incoming and outgoing CCL traffic statistics for a cluster.	bytes
Connections	Count of active connections in a cluster.	number
NAT Translations	Count of NAT translations for a cluster.	number
Distribution	Connection distribution count in the cluster for every second.	number
Packets	Packet distribution count in the cluster for every second.	number

Firewall Management Center: Troubleshooting the Cluster

You can use the **CCL Ping** tool to make sure the cluster control link is operating correctly. You can also use the following tools that are available for devices and clusters:

- Troubleshooting files—If a node fails to join the cluster, a troubleshooting file is automatically generated. You can also generate and download troubleshooting files from the **Devices** > **Device Management** > **Cluster** > **General** area. See Generate Troubleshooting Files, on page 129.

  You can also generate files from the **Device Management** page by clicking **More** (⋮) and choosing **Troubleshoot Files**.

- CLI output—From the **Devices** > **Device Management** > **Cluster** > **General** area, you can view a set of pre-defined CLI outputs that can help you troubleshoot the cluster. The following commands are automatically run for the cluster:
  - **show running-config cluster**
  - **show cluster info**
  - **show cluster info health**
  - **show cluster info transport cp**
  - **show version**
  - **show asp drop**
  - **show counters**

- **show arp**
- **show int ip brief**
- **show blocks**
- **show cpu detailed**
- **show interface** *ccl_interface*
- **ping** *ccl_ip* **size** *ccl_mtu* **repeat 2**

You can also enter any **show** command in the Command field. See View CLI Output, on page 132 for more information.

Perform a Ping on the Cluster Control Link

When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the ping fails, a notification is generated so you can fix the MTU mismatch on connecting switches and try again. This tool lets you manually ping all nodes that have already joined the cluster in case you are having cluster control link connectivity problems.

Procedure

Step 1 Choose **Devices** > **Device Management**, click the **More** (⋮) icon next to the cluster, and choose **Cluster Live Status**.

Figure 314: Cluster Status

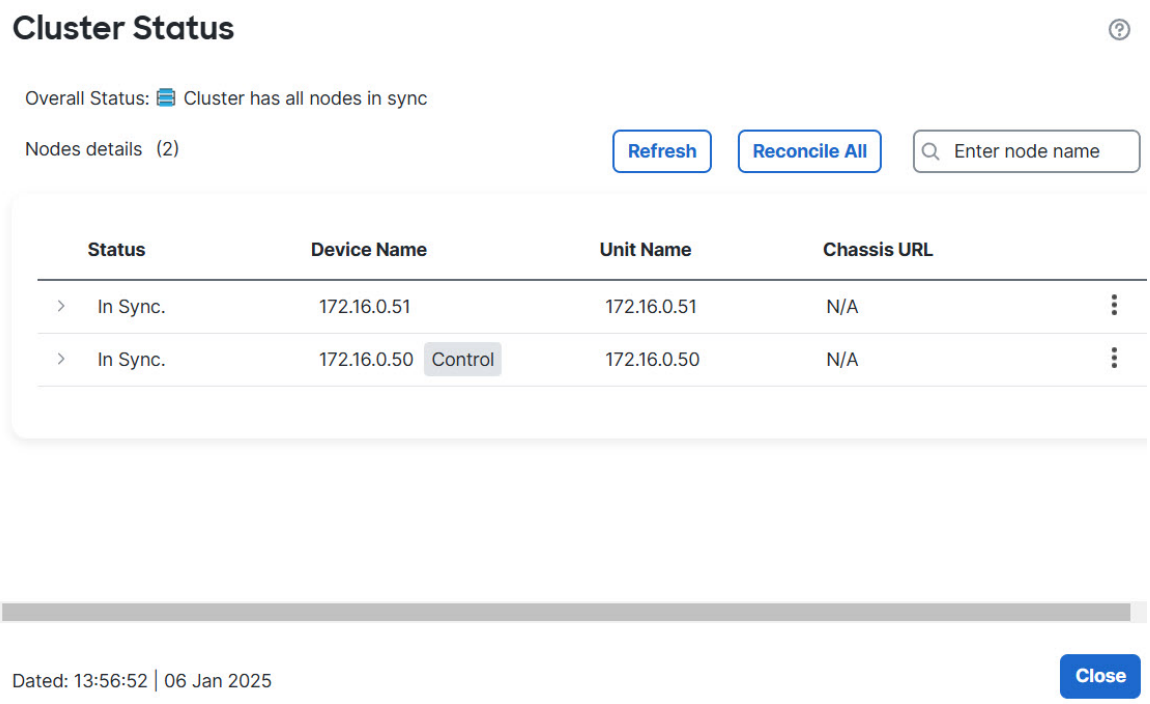

Step 2 Expand one of the nodes, and click **CCL Ping**.

Figure 315: CCL Ping

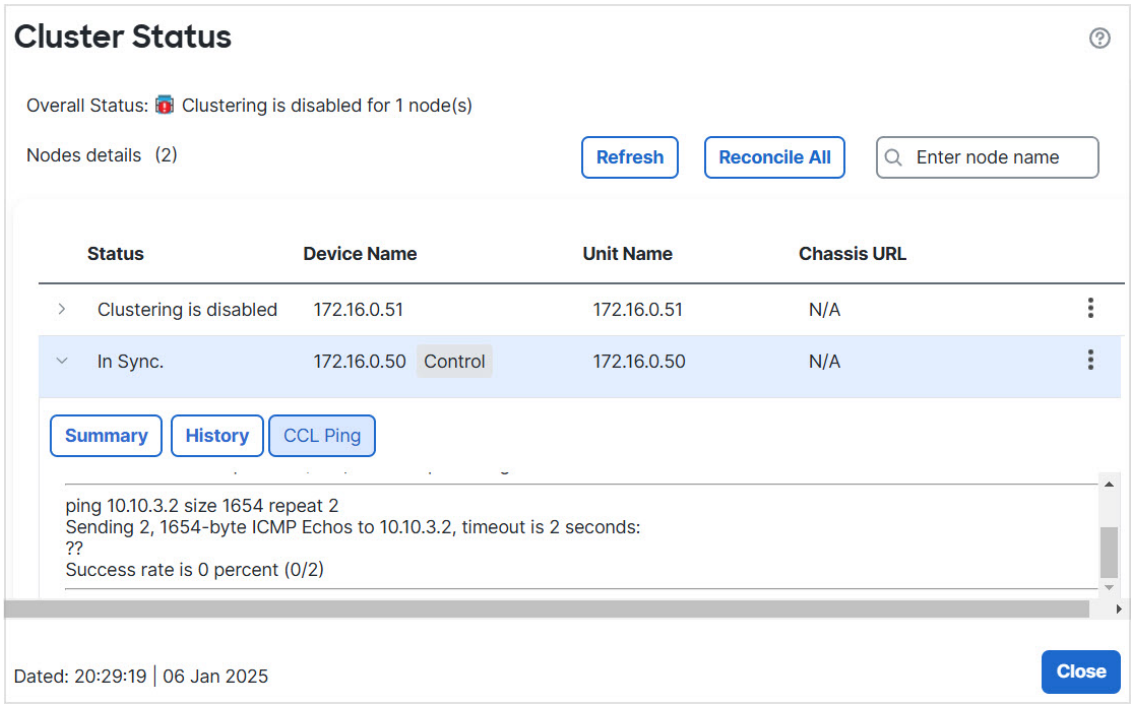

The node sends a ping on the cluster control link to every other node using a packet size that matches the maximum MTU.

Examples for Clustering

These examples include typical deployments.

Firewall on a Stick

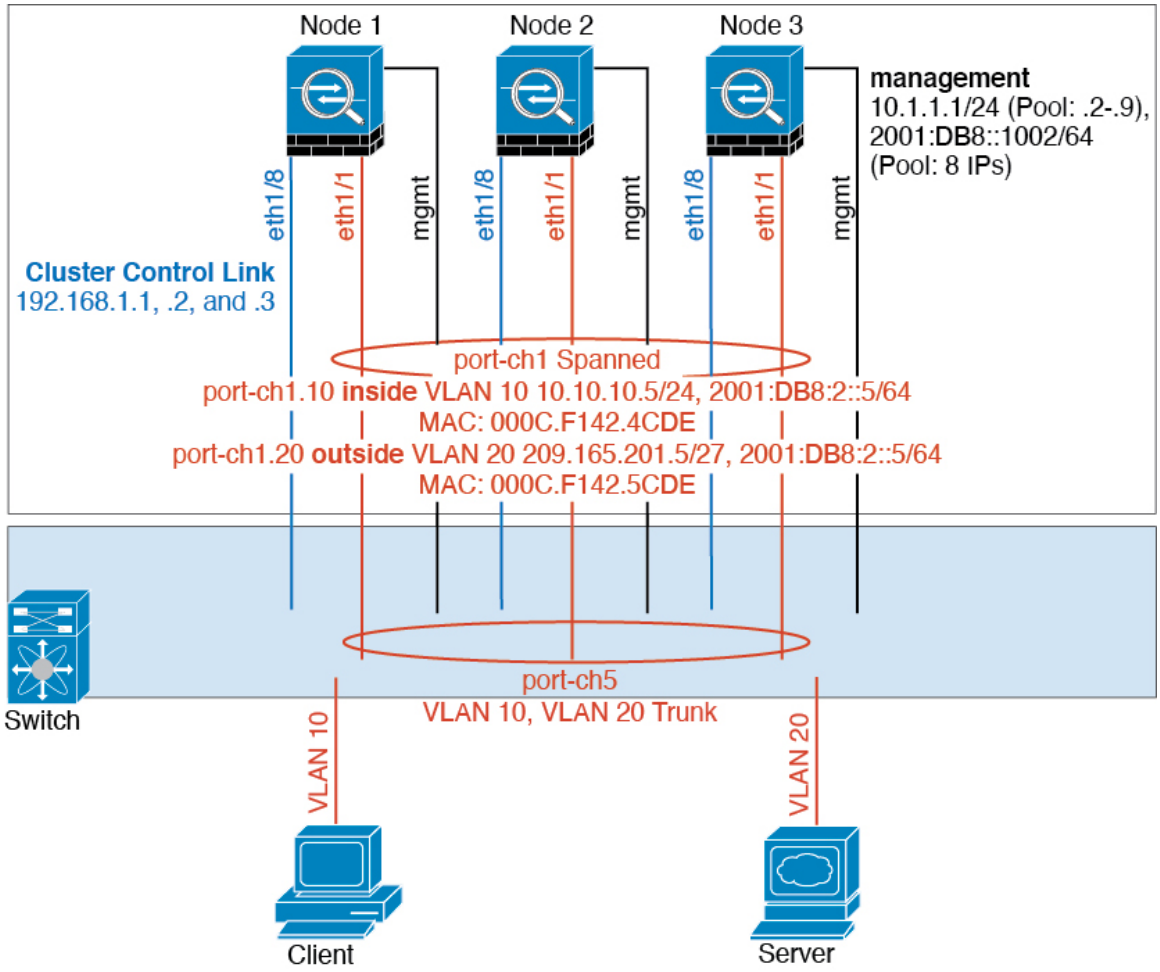

Data traffic from different security domains are associated with different VLANs, for example, VLAN 10 for the inside network and VLAN 20 for the outside network. Each has a single physical port connected to the external switch or router. Trunking is enabled so that all packets on the physical link are 802.1q encapsulated. The is the firewall between VLAN 10 and VLAN 20.

When using Spanned EtherChannels, all data links are grouped into one EtherChannel on the switch side. If the becomes unavailable, the switch will rebalance traffic between the remaining units.

Traffic Segregation

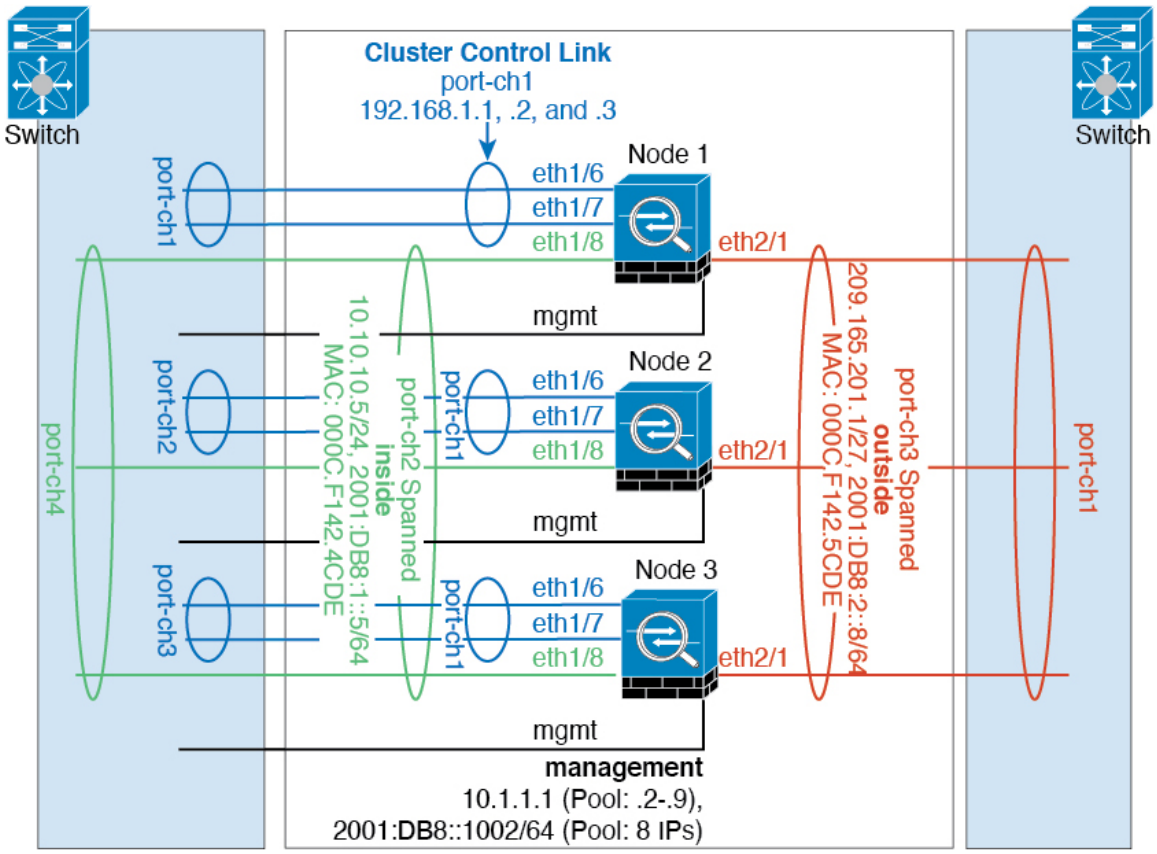

You may prefer physical separation of traffic between the inside and outside network.

As shown in the diagram above, there is one Spanned EtherChannel on the left side that connects to the inside switch, and the other on the right side to outside switch. You can also create VLAN subinterfaces on each EtherChannel if desired.

Reference for Clustering

This section includes more information about how clustering operates.

Firewall Threat Defense Features and Clustering

Some Firewall Threat Defense features are not supported with clustering, and some are only supported on the control unit. Other features might have caveats for proper usage.

Unsupported Features with Clustering

These features cannot be configured with clustering enabled, and the commands will be rejected.

Note To view FlexConfig features that are also not supported with clustering, for example WCCP inspection, see the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the Firewall Management Center GUI. See FlexConfig Policies, on page 2667.

- Remote access VPN (SSL VPN and IPsec VPN)
- DHCP client, server, and proxy. DHCP relay is supported.
- Virtual Tunnel Interfaces (VTIs)
- High Availability
- Integrated Routing and Bridging
- Firewall Management Center UCAPL/CC mode
- DHCP client, server, and proxy. DHCP relay is supported.

Centralized Features for Clustering

The following features are only supported on the control node, and are not scaled for the cluster.

Note Traffic for centralized features is forwarded from member nodes to the control node over the cluster control link.

If you use the rebalancing feature, traffic for centralized features may be rebalanced to non-control nodes before the traffic is classified as a centralized feature; if this occurs, the traffic is then sent back to the control node.

For centralized features, if the control node fails, all connections are dropped, and you have to re-establish the connections on the new control node.

Note To view FlexConfig features that are also centralized with clustering, for example RADIUS inspection, see the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the Firewall Management Center GUI. See FlexConfig Policies, on page 2667.

- The following application inspections:
  - DCERPC
  - ESMTP
  - NetBIOS
  - PPTP
  - RSH
  - SQLNET

- SUNRPC
- TFTP
- XDMCP

- Static route monitoring
- Site-to-site VPN
- IGMP multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- PIM multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- Dynamic routing

Connection Settings

Connection limits are enforced cluster-wide. Each node has an estimate of the cluster-wide counter values based on broadcast messages. Due to efficiency considerations, the configured connection limit across the cluster might not be enforced exactly at the limit number. Each node may overestimate or underestimate the cluster-wide counter value at any given time. However, the information will get updated over time in a load-balanced cluster.

Dynamic Routing and Clustering

The routing process only runs on the control unit, and routes are learned through the control unit and replicated to secondaries. If a routing packet arrives at a data unit, it is redirected to the control unit.

Figure 316: Dynamic Routing

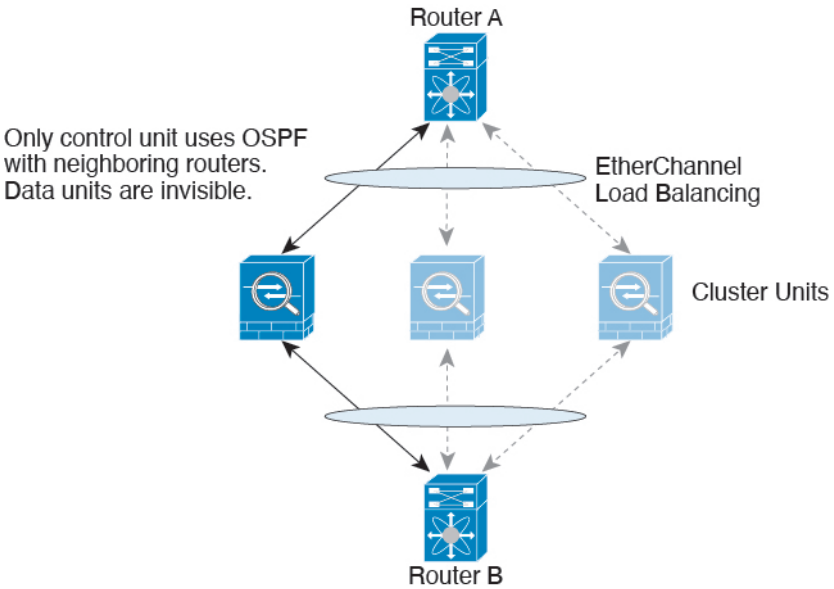

After the data units learn the routes from the control unit, each unit makes forwarding decisions independently.

The OSPF LSA database is not synchronized from the control unit to data units. If there is a control unit switchover, the neighboring router will detect a restart; the switchover is not transparent. The OSPF process picks an IP address as its router ID. Although not required, you can assign a static router ID to ensure a consistent router ID is used across the cluster. See the OSPF Non-Stop Forwarding feature to address the interruption.

FTP and Clustering

- If FTP data channel and control channel flows are owned by different cluster members, then the data channel owner will periodically send idle timeout updates to the control channel owner and update the idle timeout value. However, if the control flow owner is reloaded, and the control flow is re-hosted, the parent/child flow relationship will not longer be maintained; the control flow idle timeout will not be updated.

Multicast Routing and Clustering

The control unit handles all multicast routing packets and data packets until fast-path forwarding is established. After the connection is established, each data unit can forward multicast data packets.

NAT and Clustering

NAT can affect the overall throughput of the cluster. Inbound and outbound NAT packets can be sent to different Firewall Threat Defenses in the cluster, because the load balancing algorithm relies on IP addresses and ports, and NAT causes inbound and outbound packets to have different IP addresses and/or ports. When a packet arrives at the Firewall Threat Defense that is not the NAT owner, it is forwarded over the cluster control link to the owner, causing large amounts of traffic on the cluster control link. Note that the receiving node does not create a forwarding flow to the owner, because the NAT owner may not end up creating a connection for the packet depending on the results of security and policy checks.

If you still want to use NAT in clustering, then consider the following guidelines:

- PAT with Port Block Allocation—See the following guidelines for this feature:
  - Maximum-per-host limit is not a cluster-wide limit, and is enforced on each node individually. Thus, in a 3-node cluster with the maximum-per-host limit configured as 1, if the traffic from a host is load-balanced across all 3 nodes, then it can get allocated 3 blocks with 1 in each node.
  - Port blocks created on the backup node from the backup pools are not accounted for when enforcing the maximum-per-host limit.
  - On-the-fly PAT rule modifications, where the PAT pool is modified with a completely new range of IP addresses, will result in xlate backup creation failures for the xlate backup requests that were still in transit while the new pool became effective. This behavior is not specific to the port block allocation feature, and is a transient PAT pool issue seen only in cluster deployments where the pool is distributed and traffic is load-balanced across the cluster nodes.
  - When operating in a cluster, you cannot simply change the block allocation size. The new size is effective only after you reload each device in the cluster. To avoid having to reload each device, we recommend that you delete all block allocation rules and clear all xlates related to those rules. You can then change the block size and recreate the block allocation rules.
- NAT pool address distribution for dynamic PAT—When you configure a PAT pool, the cluster divides each IP address in the pool into port blocks. By default, each block is 512 ports, but if you configure port block allocation rules, your block setting is used instead. These blocks are distributed evenly among the

nodes in the cluster, so that each node has one or more blocks for each IP address in the PAT pool. Thus, you could have as few as one IP address in a PAT pool for a cluster, if that is sufficient for the number of PAT'ed connections you expect. Port blocks cover the 1024-65535 port range, unless you configure the option to include the reserved ports, 1-1023, on the PAT pool NAT rule.

- Reusing a PAT pool in multiple rules—To use the same PAT pool in multiple rules, you must be careful about the interface selection in the rules. You must either use specific interfaces in all rules, or "any" in all rules. You cannot mix specific interfaces and "any" across the rules, or the system might not be able to match return traffic to the right node in the cluster. Using unique PAT pools per rule is the most reliable option.
- No round-robin—Round-robin for a PAT pool is not supported with clustering.
- No extended PAT—Extended PAT is not supported with clustering.
- Dynamic NAT xlates managed by the control node—The control node maintains and replicates the xlate table to data nodes. When a data node receives a connection that requires dynamic NAT, and the xlate is not in the table, it requests the xlate from the control node. The data node owns the connection.
- Stale xlates—The xlate idle time on the connection owner does not get updated. Thus, the idle time might exceed the idle timeout. An idle timer value higher than the configured timeout with a refcnt of 0 is an indication of a stale xlate.
- No static PAT for the following inspections—
  - FTP
  - RSH
  - SQLNET
  - TFTP
  - XDMCP
  - SIP
- If you have an extremely large number of NAT rules, over ten thousand, you should enable the transactional commit model using the **asp rule-engine transactional-commit nat** command in the device CLI. Otherwise, the node might not be able to join the cluster.

SIP Inspection and Clustering

A control flow can be created on any node (due to load balancing); its child data flows must reside on the same node.

SNMP and Clustering

You should always use the Local address, and not the Main cluster IP address for SNMP polling. If the SNMP agent polls the Main cluster IP address, if a new control node is elected, the poll to the new control node will fail.

When using SNMPv3 with clustering, if you add a new cluster node after the initial cluster formation, then SNMPv3 users are not replicated to the new node. You must remove the users, and re-add them, and then redeploy your configuration to force the users to replicate to the new node.

Syslog and Clustering

- Each node in the cluster generates its own syslog messages. You can configure logging so that each node uses either the same or a different device ID in the syslog message header field. For example, the hostname configuration is replicated and shared by all nodes in the cluster. If you configure logging to use the hostname as the device ID, syslog messages generated by all nodes look as if they come from a single node. If you configure logging to use the local-node name that is assigned in the cluster bootstrap configuration as the device ID, syslog messages look as if they come from different nodes.

TLS/SSL Connections and Clustering

The decryption states of TLS/SSL connections are not synchronized, and if the connection owner fails, then the decrypted connections will be reset. New connections will need to be established to a new unit. Connections that are not decrypted (they match a do-not-decrypt rule) are not affected and are replicated correctly.

Cisco TrustSec and Clustering

Only the control node learns security group tag (SGT) information. The control node then populates the SGT to data nodes, and data nodes can make a match decision for SGT based on the security policy.

VPN and Clustering

Site-to-site VPN is a centralized feature; only the control unit supports VPN connections.

Note Remote access VPN is not supported with clustering.

VPN functionality is limited to the control unit and does not take advantage of the cluster high availability capabilities. If the control unit fails, all existing VPN connections are lost, and VPN users will see a disruption in service. When a new control unit is elected, you must reestablish the VPN connections.

When you connect a VPN tunnel to a Spanned interface address, connections are automatically forwarded to the control unit.

VPN-related keys and certificates are replicated to all units.

Performance Scaling Factor

When you combine multiple units into a cluster, you can expect the total cluster performance to be approximately 80% of the maximum combined throughput.

For example, for TCP throughput, the Firepower 9300 with 3 SM-40 modules can handle approximately 135 Gbps of real world firewall traffic when running alone. For 2 chassis, the maximum combined throughput will be approximately 80% of 270 Gbps (2 chassis x 135 Gbps): 216 Gbps.

Control Unit Election

Members of the cluster communicate over the cluster control link to elect a control unit as follows:

1. When you deploy the cluster, each unit broadcasts an election request every 3 seconds.

2. Any other units with a higher priority respond to the election request; the priority is set when you deploy the cluster and is not configurable.
3. If after 45 seconds, a unit does not receive a response from another unit with a higher priority, then it becomes the control unit.

Note If multiple units tie for the highest priority, the cluster unit name and then the serial number is used to determine the control unit.

4. If a unit later joins the cluster with a higher priority, it does not automatically become the control unit; the existing control unit always remains as the control unit unless it stops responding, at which point a new control unit is elected.
5. In a "split brain" scenario when there are temporarily multiple control units, then the unit with highest priority retains the role while the other units return to data unit roles.

Note You can manually force a unit to become the control unit. For centralized features, if you force a control unit change, then all connections are dropped, and you have to re-establish the connections on the new control unit.

High Availability Within the Cluster

Clustering provides high availability by monitoring chassis, unit, and interface health and by replicating connection states between units.

Chassis-Application Monitoring

Chassis-application health monitoring is always enabled. The Firepower 4100/9300 chassis supervisor checks the Firewall Threat Defense application periodically (every second). If the Firewall Threat Defense device is up and cannot communicate with the Firepower 4100/9300 chassis supervisor for 3 seconds, the Firewall Threat Defense device generates a syslog message and leaves the cluster.

If the Firepower 4100/9300 chassis supervisor cannot communicate with the application after 45 seconds, it reloads the Firewall Threat Defense device. If the Firewall Threat Defense device cannot communicate with the supervisor, it removes itself from the cluster.

Unit Health Monitoring

Each unit periodically sends a broadcast keepaliveheartbeat packet over the cluster control link. If the control node does not receive any keepaliveheartbeat packets or other packets from a data node within the configurable timeout period, then the control node removes the data node from the cluster. If the data nodes do not receive packets from the control node, then a new control node is elected from the remaining node.

If nodes cannot reach each other over the cluster control link because of a network failure and not because a node has actually failed, then the cluster may go into a "split brain" scenario where isolated data nodes will elect their own control nodes. For example, if a router fails between two cluster locations, then the original control node at location 1 will remove the location 2 data nodes from the cluster. Meanwhile, the nodes at location 2 will elect their own control node and form their own cluster. Note that asymmetric traffic may fail

in this scenario. After the cluster control link is restored, then the control node that has the higher priority will keep the control node's role. See Control Unit Election, on page 758 for more information.

Interface Monitoring

Each node monitors the link status of all hardware interfaces in use, and reports status changes to the control node. For clustering on multiple chassis, Spanned EtherChannels use the cluster Link Aggregation Control Protocol (cLACP). Each chassis monitors the link status and the cLACP protocol messages to determine if the port is still active in the EtherChannel, and informs the Firewall Threat Defense application if the interface is down. When you enable health monitoring, all physical interfaces are monitored by default (including the main EtherChannel for EtherChannel interfaces). Only named interfaces that are in an Up state can be monitored. For example, all member ports of an EtherChannel must fail before a *named* EtherChannel is removed from the cluster. You can optionally disable monitoring per interface.

If a monitored interface fails on a particular node, but it is active on other nodes, then the node is removed from the cluster. The amount of time before the Firewall Threat Defense device removes a node from the cluster depends on whether the node is an established member or is joining the cluster. The Firewall Threat Defense device does not monitor interfaces for the first 90 seconds that a node joins the cluster. Interface status changes during this time will not cause the Firewall Threat Defense device to be removed from the cluster. For an established member, the node is removed after 500 ms.

For clustering on multiple chassis, if you add or delete an EtherChannel from the cluster, interface health-monitoring is suspended for 95 seconds to ensure that you have time to make the changes on each chassis.

Decorator Application Monitoring

When you install a decorator application on an interface, such as the Radware DefensePro application, then both the Firewall Threat Defense device and the decorator application must be operational to remain in the cluster. The unit does not join the cluster until both applications are operational. Once in the cluster, the unit monitors the decorator application health every 3 seconds. If the decorator application is down, the unit is removed from the cluster.

Status After Failure

When a node in the cluster fails, the connections hosted by that node are seamlessly transferred to other nodes; state information for traffic flows is shared over the control node's cluster control link.

If the control node fails, then another member of the cluster with the highest priority (lowest number) becomes the control node.

The Firewall Threat Defense automatically tries to rejoin the cluster, depending on the failure event.

Note When the Firewall Threat Defense becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the Management interface can send and receive traffic.

Rejoining the Cluster

After a cluster member is removed from the cluster, how it can rejoin the cluster depends on why it was removed:

- Failed cluster control link when initially joining—After you resolve the problem with the cluster control link, you must manually rejoin the cluster by re-enabling clustering.
- Failed cluster control link after joining the cluster—The Firewall Threat Defense automatically tries to rejoin every 5 minutes, indefinitely.
- Failed data interface—The Firewall Threat Defense automatically tries to rejoin at 5 minutes, then at 10 minutes, and finally at 20 minutes. If the join is not successful after 20 minutes, then the Firewall Threat Defense application disables clustering. After you resolve the problem with the data interface, you have to manually enable clustering.
- Failed node—If the node was removed from the cluster because of a node health check failure, then rejoining the cluster depends on the source of the failure. For example, a temporary power failure means the node will rejoin the cluster when it starts up again as long as the cluster control link is up. The Firewall Threat Defense application attempts to rejoin the cluster every 5 seconds.
- Internal error—Internal failures include: application sync timeout; inconsistent application statuses; and so on.
- Failed configuration deployment—If you deploy a new configuration from Firewall Management Center, and the deployment fails on some cluster members but succeeds on others, then the nodes that failed are removed from the cluster. You must manually rejoin the cluster by re-enabling clustering. If the deployment fails on the control node, then the deployment is rolled back, and no members are removed. If the deployment fails on all data nodes, then the deployment is rolled back, and no members are removed.
- Failed Chassis-Application Communication—When the Firewall Threat Defense application detects that the chassis-application health has recovered, it tries to rejoin the cluster automatically.

Data Path Connection State Replication

Every connection has one owner and at least one backup owner in the cluster. The backup owner does not take over the connection in the event of a failure; instead, it stores TCP/UDP state information, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner is usually also the director.

Some traffic requires state information above the TCP or UDP layer. See the following table for clustering support or lack of support for this kind of traffic.

Table 48: Features Replicated Across the Cluster

Traffic	State Support	Notes
Up time	Yes	Keeps track of the system up time.
ARP Table	Yes	—
MAC address table	Yes	—
User Identity	Yes	—
IPv6 Neighbor database	Yes	—
Dynamic routing	Yes	—
SNMP Engine ID	**No**	—

How the Cluster Manages Connections

Connections can be load-balanced to multiple nodes of the cluster. Connection roles determine how connections are handled in both normal operation and in a high availability situation.

Connection Roles

See the following roles defined for each connection:

- Owner—Usually, the node that initially receives the connection. The owner maintains the TCP state and processes packets. A connection has only one owner. If the original owner fails, then when new nodes receive packets from the connection, the director chooses a new owner from those nodes.
- Backup owner—The node that stores TCP/UDP state information received from the owner, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner does not take over the connection in the event of a failure. If the owner becomes unavailable, then the first node to receive packets from the connection (based on load balancing) contacts the backup owner for the relevant state information so it can become the new owner.

  As long as the director (see below) is not the same node as the owner, then the director is also the backup owner. If the owner chooses itself as the director, then a separate backup owner is chosen.

  For clustering on the Firepower 9300, which can include up to 3 cluster nodes in one chassis, if the backup owner is on the same chassis as the owner, then an additional backup owner will be chosen from another chassis to protect flows from a chassis failure.
- Director—The node that handles owner lookup requests from forwarders. When the owner receives a new connection, it chooses a director based on a hash of the source/destination IP address and ports (see below for ICMP hash details), and sends a message to the director to register the new connection. If packets arrive at any node other than the owner, the node queries the director about which node is the owner so it can forward the packets. A connection has only one director. If a director fails, the owner chooses a new director.

  As long as the director is not the same node as the owner, then the director is also the backup owner (see above). If the owner chooses itself as the director, then a separate backup owner is chosen.

  ICMP/ICMPv6 hash details:

  - For Echo packets, the source port is the ICMP identifier, and the destination port is 0.
  - For Reply packets, the source port is 0, and the destination port is the ICMP identifier.
  - For other packets, both source and destination ports are 0.
- Forwarder—A node that forwards packets to the owner. If a forwarder receives a packet for a connection it does not own, it queries the director for the owner, and then establishes a flow to the owner for any other packets it receives for this connection. The director can also be a forwarder. Note that if a forwarder receives the SYN-ACK packet, it can derive the owner directly from a SYN cookie in the packet, so it does not need to query the director. (If you disable TCP sequence randomization, the SYN cookie is not used; a query to the director is required.) For short-lived flows such as DNS and ICMP, instead of querying, the forwarder immediately sends the packet to the director, which then sends them to the owner. A connection can have multiple forwarders; the most efficient throughput is achieved by a good load-balancing method where there are no forwarders and all packets of a connection are received by the owner.

Note We do not recommend disabling TCP sequence randomization when using clustering. There is a small chance that some TCP sessions won't be established, because the SYN/ACK packet might be dropped.

- Fragment Owner—For fragmented packets, cluster nodes that receive a fragment determine a fragment owner using a hash of the fragment source IP address, destination IP address, and the packet ID. All fragments are then forwarded to the fragment owner over the cluster control link. Fragments may be load-balanced to different cluster nodes, because only the first fragment includes the 5-tuple used in the switch load balance hash. Other fragments do not contain the source and destination ports and may be load-balanced to other cluster nodes. The fragment owner temporarily reassembles the packet so it can determine the director based on a hash of the source/destination IP address and ports. If it is a new connection, the fragment owner will register to be the connection owner. If it is an existing connection, the fragment owner forwards all fragments to the provided connection owner over the cluster control link. The connection owner will then reassemble all fragments.

Port Address Translation Connections

New Connection Ownership

When a new connection is directed to a node of the cluster via load balancing, that node owns both directions of the connection. If any connection packets arrive at a different node, they are forwarded to the owner node over the cluster control link. If a reverse flow arrives at a different node, it is redirected back to the original node.

Sample Data Flow for TCP

The following example shows the establishment of a new connection.

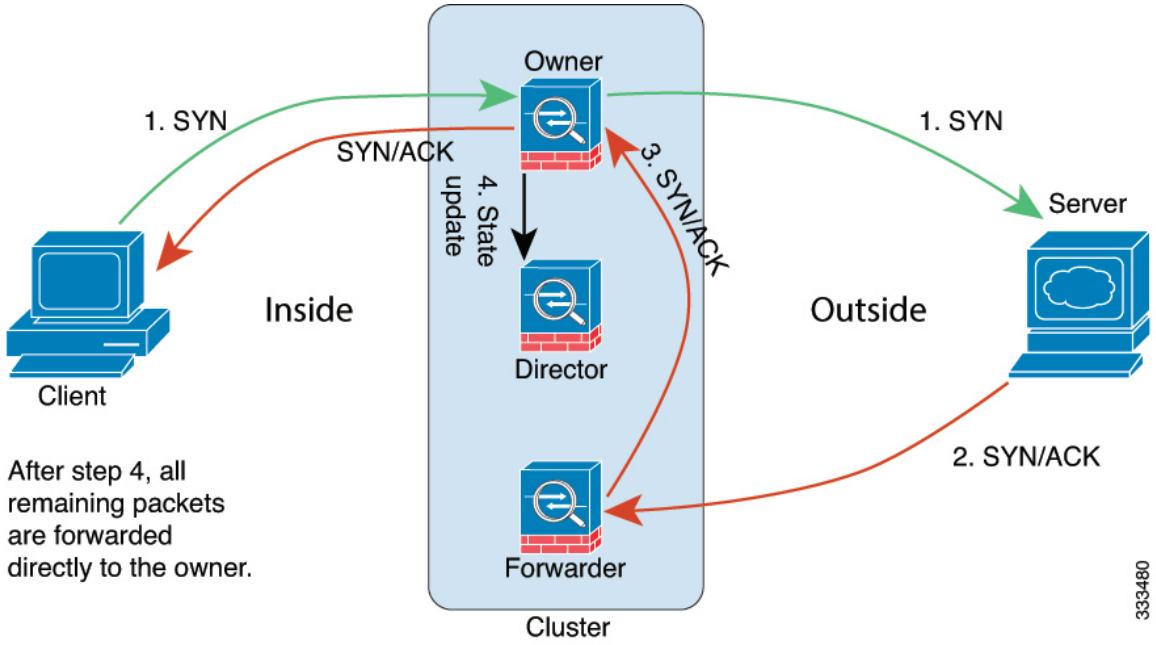

1. The SYN packet originates from the client and is delivered to one Firewall Threat Defense (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.
2. The SYN-ACK packet originates from the server and is delivered to a different Firewall Threat Defense (based on the load balancing method). This Firewall Threat Defense is the forwarder.
3. Because the forwarder does not own the connection, it decodes owner information from the SYN cookie, creates a forwarding flow to the owner, and forwards the SYN-ACK to the owner.
4. The owner sends a state update to the director, and forwards the SYN-ACK to the client.
5. The director receives the state update from the owner, creates a flow to the owner, and records the TCP state information as well as the owner. The director acts as the backup owner for the connection.
6. Any subsequent packets delivered to the forwarder will be forwarded to the owner.
7. If packets are delivered to any additional nodes, it will query the director for the owner and establish a flow.
8. Any state change for the flow results in a state update from the owner to the director.

Sample Data Flow for ICMP and UDP

The following example shows the establishment of a new connection.

1. ***Figure 317: ICMP and UDP Data Flow***

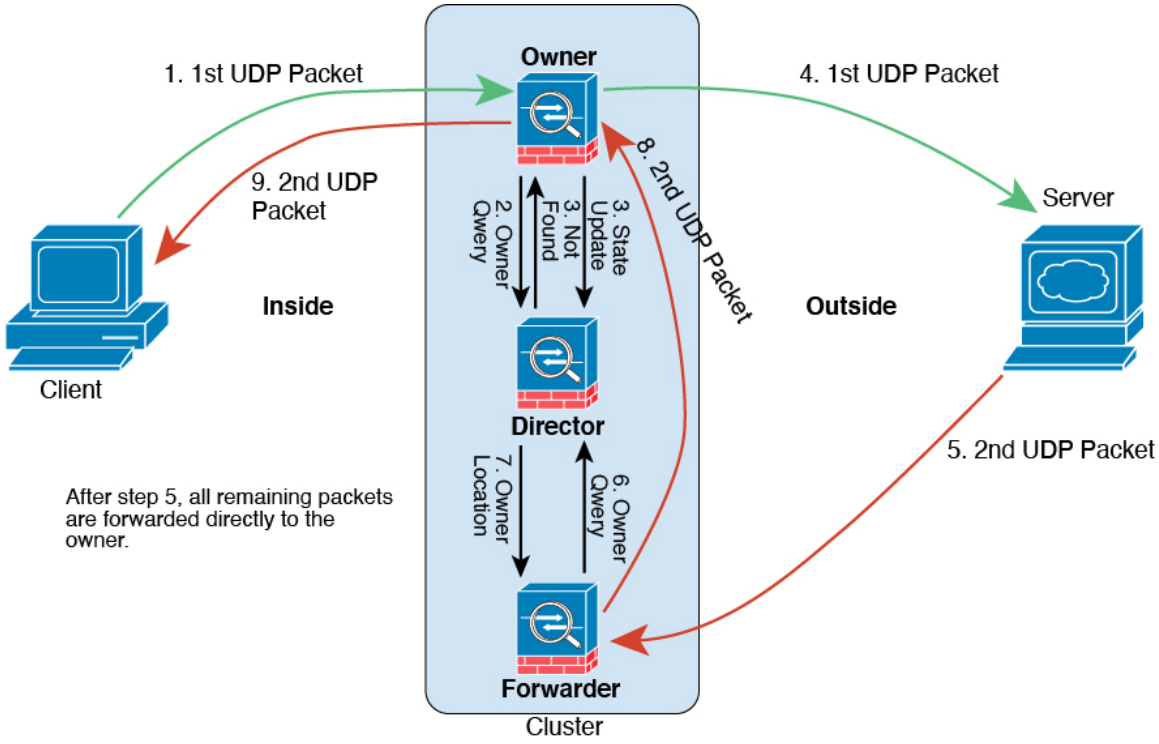

The first UDP packet originates from the client and is delivered to one Firewall Threat Defense (based on the load balancing method).

2. The node that received the first packet queries the director node that is chosen based on a hash of the source/destination IP address and ports.
3. The director finds no existing flow, creates a director flow and forwards the packet back to the previous node. In other words, the director has elected an owner for this flow.
4. The owner creates the flow, sends a state update to the director, and forwards the packet to the server.
5. The second UDP packet originates from the server and is delivered to the forwarder.
6. The forwarder queries the director for ownership information. For short-lived flows such as DNS, instead of querying, the forwarder immediately sends the packet to the director, which then sends it to the owner.
7. The director replies to the forwarder with ownership information.
8. The forwarder creates a forwarding flow to record owner information and forwards the packet to the owner.
9. The owner forwards the packet to the client.

History for Clustering

Table 49:

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
MTU ping test on cluster node join	7.6.0	7.6.0	When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the ping fails, a notification is generated so you can fix the MTU mismatch on connecting switches and try again.
Cluster control link ping tool.	7.2.6/7.4.1	Any	You can check to make sure all the cluster nodes can reach each other over the cluster control link by performing a ping. One major cause for the failure of a node to join the cluster is an incorrect cluster control link configuration; for example, the cluster control link MTU may be set higher than the connecting switch MTUs. New/modified screens: **Devices** > **Device Management** > **More** > **Cluster Live Status.** Version restrictions: Not supported with Firewall Management Center Version 7.3.x or 7.4.0.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Troubleshooting file generation and download available from Device and Cluster pages.	7.4.1	7.4.1	You can generate and download troubleshooting files for each device on the Device page and also for all cluster nodes on the Cluster page. For a cluster, you can download all files as a single compressed file. You can also include cluster logs for the cluster for cluster nodes. You can alternatively trigger file generation from the **Devices** > **Device Management** > **More** > **Troubleshoot Files** menu. New/modified screens: • **Devices** > **Device Management** > **Device** > **General** • **Devices** > **Device Management** > **Cluster** > **General**
View CLI output for a device or device cluster.	7.4.1	Any	You can view a set of pre-defined CLI outputs that can help you troubleshoot the device or cluster. You can also enter any **show** command and see the output. New/modified screens: **Devices** > **Device Management** > **Cluster** > **General**
Cluster health monitor settings.	7.3.0	Any	You can now edit cluster health monitor settings. New/Modified screens: **Devices** > **Device Management** > **Cluster** > **Cluster Health Monitor Settings** **Note** If you previously configured these settings using FlexConfig, be sure to remove the FlexConfig configuration before you deploy. Otherwise the FlexConfig configuration will overwrite the management center configuration.
Cluster health monitor dashboard.	7.3.0	Any	You can now view cluster health on the cluster health monitor dashboard. New/Modified screens: **System** > **Health** > **Monitor**
Support for 16-node clusters.	7.2.0	7.2.0	You can now configure 16 node clusters for the Firepower 4100/9300. Previously, the maximum was 6 units. New/Modified screens: none. Supported platforms: Firepower 4100/9300
Cluster deployment for firewall changes completes faster.	7.1.0	7.1.0	Cluster deployment for firewall changes now completes faster. New/Modified screens: none.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Improved PAT port block allocation for clustering.	7.0.0	7.0.0	The improved PAT port block allocation ensures that the control unit keeps ports in reserve for joining nodes, and proactively reclaims unused ports. To best optimize the allocation, you can set the maximum nodes you plan to have in the cluster using the **cluster-member-limit** command using FlexConfig. The control unit can then allocate port blocks to the planned number of nodes, and it will not have to reserve ports for extra nodes you don't plan to use. The default is 16 nodes. You can also monitor syslog 747046 to ensure that there are enough ports available for a new node. New/Modified commands: **cluster-member-limit** (FlexConfig), **show nat pool cluster** [**summary**], **show nat pool ip detail**
Cluster deployment for Snort changes completes faster, and fails faster when there is an event.	6.7.0	6.7.0	Cluster deployment for Snort changes now completes faster. Also, when a cluster has an event that causes a Firewall Management Center deployment to fail, the failure now occurs more quickly. New/Modified screens: none.
Improved cluster management.	6.7.0	6.7.0	Firewall Management Center has improved cluster management functionality that formerly you could only accomplish using the CLI, including: • Enable and disable cluster units • Show cluster status from the Device Management page, including History and Summary per unit • Change the role to the control unit New/Modified screens: • **Devices** > **Device Management** > **More** menu • **Devices** > **Device Management** > **Cluster** > **General** area **> Cluster Live Status** link **Cluster Status** Supported platforms: Firepower 4100/9300

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Multi-instance clustering.	6.6.0	6.6.0	You can now create a cluster using container instances. On the Firepower 9300, you must include one container instance on each module in the cluster. You cannot add more than one container instance to the cluster per security engine/module. We recommend that you use the same security module or chassis model for each cluster instance. However, you can mix and match container instances on different Firepower 9300 security module types or Firepower 4100 models in the same cluster if required. You cannot mix Firepower 9300 and 4100 instances in the same cluster. New/Modified FXOS commands: **set port-type cluster** New/modified Firepower Chassis Manager screens: • **Logical Devices** > **Add Cluster** • **Interfaces** > **All Interfaces** > **Add New** drop-down menu**Subinterface Type field** Supported platforms: Firewall Threat Defense on the Firepower 4100/9300
Configuration sync to data units in parallel.	6.6.0	6.6.0	The control unit now syncs configuration changes with data units in parallel by default. Formerly, synching occurred sequentially. New/Modified screens: none.
Messages for cluster join failure or eviction added to **show cluster history**.	6.6.0	6.6.0	New messages were added to the **show cluster history** command for when a cluster unit either fails to join the cluster or leaves the cluster. New/Modified commands: **show cluster history** New/Modified screens: none.
Initiator and responder information for Dead Connection Detection (DCD), and DCD support in a cluster.	6.5.0	6.5.0	If you enable Dead Connection Detection (DCD), you can use the **show conn detail** command to get information about the initiator and responder. Dead Connection Detection allows you to maintain an inactive connection, and the **show conn** output tells you how often the endpoints have been probed. In addition, DCD is now supported in a cluster. New/Modified commands: **show conn** (output only). Supported platforms: Firewall Threat Defense on the Firepower 4100/9300

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Adding clusters is easier.	6.3.0	6.3.0	You can now add any unit of a cluster to the Firewall Management Center, and the other cluster units are detected automatically. Formerly, you had to add each cluster unit as a separate device, and then group them into a cluster. Adding a cluster unit is also now automatic. Note that you must delete a unit manually. New/Modified screens: **Devices** > **Device Management** > **Add** drop-down menu **Device** > **Add Device** dialog box **Devices** > **Device Management** > **Cluster** tab**General** area**Cluster Registration Status** **Current Cluster Summary** link **Cluster Status** dialog box Supported platforms: Firewall Threat Defense on the Firepower 4100/9300
Support for site-to-site VPN with clustering as a centralized feature.	6.2.3.3	6.2.3.3	You can now configure site-to-site VPN with clustering. Site-to-site VPN is a centralized feature; only the control unit supports VPN connections. Supported platforms: Firewall Threat Defense on the Firepower 4100/9300
Automatically rejoin the cluster after an internal failure.	6.2.3	6.2.3	Formerly, many internal error conditions caused a cluster unit to be removed from the cluster, and you were required to manually rejoin the cluster after resolving the issue. Now, a unit will attempt to rejoin the cluster automatically at the following intervals: 5 minutes, 10 minutes, and then 20 minutes. Internal failures include: application sync timeout; inconsistent application statuses; and so on. New/Modified command: **show cluster info auto-join** No modified screens. Supported platforms: Firewall Threat Defense on the Firepower 4100/9300
Clustering on multiple chassis for 6 modules; Firepower 4100 support.	6.2.0	6.2.0	With FXOS 2.1.1, you can now enable clustering on multiple chassis of the Firepower 9300 and 4100. For the Firepower 9300, you can include up to 6 modules. For example, you can use 1 module in 6 chassis, or 2 modules in 3 chassis, or any combination that provides a maximum of 6 modules. For the Firepower 4100, you can include up to 6 chassis. **Note** Inter-site clustering is also supported. However, customizations to enhance redundancy and stability, such as site-specific MAC and IP addresses, director localization, site redundancy, and cluster flow mobility, are only configurable using the FlexConfig feature. No modified screens. Supported platforms: Firewall Threat Defense on the Firepower 4100/9300

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Clustering on multiple modules with one Firepower 9300 chassis.	6.0.1	6.0.1	You can cluster up to 3 security modules within the Firepower 9300 chassis. All modules in the chassis must belong to the cluster. New/Modified screens: **Devices** > **Device Management** > **Add** > **Add Cluster** **Devices** > **Device Management** > **Cluster** Supported platforms: Firewall Threat Defense on the Firepower 9300

PART III

Interfaces and Device Settings

- Interface Overview, on page 773
- Regular Firewall Interfaces, on page 817
- Inline Sets and Passive Interfaces, on page 893
- DHCP and DDNS, on page 909
- SNMP for the Firepower 1000, on page 927
- Quality of Service, on page 933
- Platform Settings, on page 945
- Network Address Translation, on page 1015
- Alarms for the Cisco ISA 3000, on page 1135

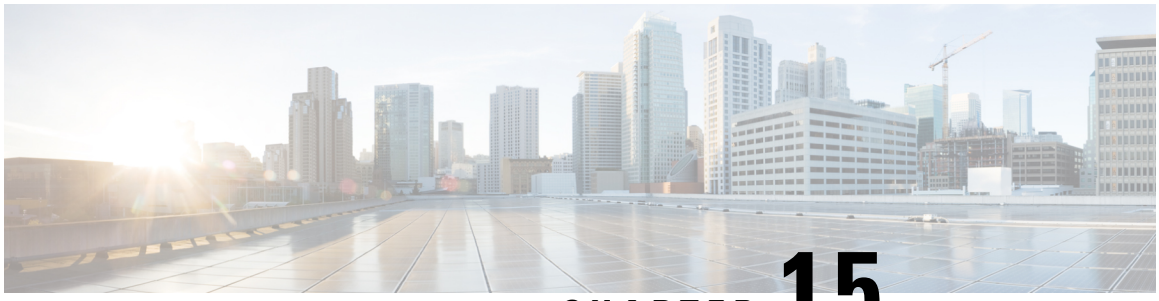

CHAPTER 15

Interface Overview

The Firewall Threat Defense device includes data interfaces that you can configure in different modes, as well as a management interface.

- Management Interface, on page 773
- Interface Mode and Types, on page 774
- Security Zones and Interface Groups, on page 775
- Auto-MDI/MDIX Feature, on page 776
- Redundant Interfaces (Deprecated), on page 777
- Default Settings for Interfaces, on page 777
- Create Security Zone and Interface Group Objects, on page 777
- Enable the Physical Interface and Configure Ethernet Settings, on page 778
- Configure EtherChannel Interfaces, on page 781
- Sync Interface Changes with the Firewall Management Center, on page 789
- Manage the Network Module for the Secure Firewall 3100/4200, on page 790
- Merge the Management and Diagnostic Interfaces, on page 805
- History for Interfaces, on page 812

Management Interface

In Version 7.3 and earlier, the physical management interface is shared between the Diagnostic logical interface and the Management logical interface. In Version 7.4 and later, the Diagnostic interface is merged with Management for a simplified user experience.

Management Interface

The Management interface is separate from the other interfaces on the device. It is used to set up and register the device to the Firewall Management Center. It uses its own IP address and static routing. You can configure its settings at the CLI using the **configure network** command. You can also view its status on the **Devices** > **Device Management** page. If you change the IP address at the CLI after you add it to the Firewall Management Center, you can match the IP address in the Secure Firewall Management Center in the **Devices** > **Device Management** area.

You can alternatively manage the Firewall Threat Defense using a data interface instead of the Management interface.

Diagnostic Interface

For new devices using 7.4 and later, you cannot use the legacy Diagnostic interface. Only the merged Management interface is available.

If you upgraded to 7.4 or later, and you did not have any configuration for the Diagnostic interface, then the interfaces will merge automatically.

If you upgraded to 7.4 or later, and you have configuration for the Diagnostic interface, then you have the choice to merge the interfaces manually, or you can continue to use the separate Diagnostic interface. Note that support for the Diagnostic interface will be removed in a later release, so you should plan to merge the interfaces as soon as possible. To manually merge the Management and Diagnostic interfaces, see Merge the Management and Diagnostic Interfaces, on page 805. Configurations that prevent an automatic merge include the following:

- A data interface named "management"—This name is reserved for use with the merged Management interface.
- IP Address on Diagnostic
- DNS enabled on Diagnostic
- Syslog, SNMP, RADIUS or AD (for remote access VPN) source interface is Diagnostic
- RADIUS or AD (for remote access VPN) with no source interface specified, and there is at least one interface configured as management-only (including Diagnostic)—The default route lookup for these services has changed from the management-only routing table to the data routing table, with no fallback to management. Therefore, you cannot use a management-only interface other than Management.
- Static routes on Diagnostic
- Dynamic routing on Diagnostic
- HTTP server on Diagnostic
- ICMP on Diagnostic
- DDNS for Diagnostic
- FlexConfig using Diagnostic

For more information about how the legacy Diagnostic interface operates, see the 7.3 version of this guide.

Interface Mode and Types

You can deploy Firewall Threat Defense interfaces in two modes: Regular firewall mode and IPS-only mode. You can include both firewall and IPS-only interfaces on the same device.

Regular Firewall Mode

Firewall mode interfaces subject traffic to firewall functions such as maintaining flows, tracking flow states at both IP and TCP layers, IP defragmentation, and TCP normalization. You can also optionally configure IPS functions for this traffic according to your security policy.

The types of firewall interfaces you can configure depends on the firewall mode set for the device: routed or transparent mode. See Transparent or Routed Firewall Mode, on page 285 for more information.

- Routed mode interfaces (routed firewall mode only)—Each interface that you want to route between is on a different subnet.
- Bridge group interfaces (routed and transparent firewall mode)—You can group together multiple interfaces on a network, and the Firewall Threat Defense device uses bridging techniques to pass traffic between the interfaces. Each bridge group includes a Bridge Virtual Interface (BVI) to which you assign an IP address on the network. In routed mode, the Firewall Threat Defense device routes between BVIs and regular routed interfaces. In transparent mode, each bridge group is separate and cannot communicate with each other.

IPS-Only Mode

You can configure your device in either a passive or inline IPS deployment. In a passive deployment, you deploy the system out-of-band from the flow of network traffic. In an inline deployment, you configure the system transparently on a network segment by binding two interfaces together.

Security Zones and Interface Groups

Each interface can be assigned to a *security zone* and/or *interface group*. You then apply your security policy based on zones or groups. For example, you can assign the "inside" interface on one or more devices to the "inside" zone; and the "outside" interfaces to the "outside" zone. You can then configure your access control policy to enable traffic to go from the inside zone to the outside zone for every device using the same zones.

To view the interfaces that belong to each object, choose **Objects** > **Object Management** and click **Interface**. This page lists the security zones and interface groups configured on your managed devices. You can expand each interface object to view the type of interfaces in each interface object.

Note Policies that apply to **any** zone (a global policy) apply to interfaces in zones as well as any interfaces that are not assigned to a zone.

Note The Management interface does not belong to a zone or interface group.

Security Zones Vs. Interface Groups

There are two types of interface objects:

- Security zones—An interface can belong to only one security zone.
- Interface groups—An interface can belong to multiple interface groups (and to one security zone).

  You can use interface groups in NAT policies, prefilter policies, and QoS policies, as well as features that let you specify the interface name directly, such as Syslog servers or DNS servers.

Some policies only support security zones, while other policies support zones and groups. Unless you need the functionality an interface group provides, you should default to using security zones because security zones are supported for all features.

You cannot change an existing security zone to an interface group or vice-versa; instead you must create a new interface object.

Note Although tunnel zones are not interface objects, you can use them in place of security zones in certain configurations; see Tunnel Zones and Prefiltering, on page 1928.

Interface Object Types

See the following interface object types:

- Passive—For IPS-only passive or ERSPAN interfaces.
- Inline—For IPS-only inline set interfaces.
- Switched—For regular firewall bridge group interfaces.
- Routed—For regular firewall routed interfaces.
- ASA—(Security zones only) For legacy ASA FirePOWER device interfaces.
- Management—(Interface groups only) For management-only interfaces.
- Loopback—(Interface groups only) For loopback interfaces.

All interfaces in an interface object must be of the same type. After you create an interface object, you cannot change the type of interfaces it contains.

Interface Names

Note that the interface (or zone name) itself does not provide any default behavior in regards to the security policy. We recommend using names that are self-describing to avoid mistakes in future configuration. A good name signifies a logical segment or traffic specification, for example:

- Names of internal interfaces—InsideV110, InsideV160, InsideV195
- Names of DMZ interfaces—DMZV11, DMZV12, DMZV-TEST
- Names of external interfaces—Outside-ASN78, Outside-ASN91

Auto-MDI/MDIX Feature

For RJ-45 interfaces, the default auto-negotiation setting also includes the Auto-MDI/MDIX feature. Auto-MDI/MDIX eliminates the need for crossover cabling by performing an internal crossover when a straight cable is detected during the auto-negotiation phase. Either the speed or duplex must be set to auto-negotiate to enable Auto-MDI/MDIX for the interface. If you explicitly set both the speed and duplex to a fixed value, thus disabling auto-negotiation for both settings, then Auto-MDI/MDIX is also disabled. For Gigabit Ethernet, when the speed and duplex are set to 1000 and full, then the interface always auto-negotiates; therefore Auto-MDI/MDIX is always enabled and you cannot disable it.

Redundant Interfaces (Deprecated)

Redundant interfaces were supported for the ASA 5500-X platforms only. We don't recommend configuring them for other platforms.

Default Settings for Interfaces

This section lists default settings for interfaces.

Default State of Interfaces

The default state of an interface depends on the type.

- Physical interfaces—Disabled. The exception is the Management interface that is enabled for initial setup. Physical interfaces includes switch ports.
- VLAN subinterfaces—Enabled. However, for traffic to pass through the subinterface, the physical interface must also be enabled.
- EtherChannel port-channel interfaces (ISA 3000)—Enabled. However, for traffic to pass through the EtherChannel, the channel group physical interfaces must also be enabled.
- EtherChannel port-channel interfaces (Firepower and Secure Firewall models)—Disabled.

Note For the Firepower 4100/9300, you can administratively enable and disable interfaces in both the chassis and in the Firewall Management Center. For an interface to be operational, the interface must be enabled in both operating systems. Because the interface state is controlled independently, you may have a mismatch between the chassis and Firewall Management Center.

Default Speed and Duplex

By default, the speed and duplex for copper (RJ-45) interfaces are set to auto-negotiate.

By default, the speed and duplex for fiber (SFP) interfaces are set to the maximum speed, with auto-negotiation enabled.

For the Secure Firewall 3100/4200, the speed is set to detect the installed SFP speed.

Create Security Zone and Interface Group Objects

Add security zones and interface groups to which you can assign device interfaces.

Tip You can create empty interface objects and add interfaces to them later. To add an interface, the interface must have a name. You can also create security zones (but not interface groups) while configuring interfaces.

Before you begin

Understand the usage requirements and restrictions for each type of interface object. See Security Zones and Interface Groups, on page 775.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Interface** from the list of object types.

Step 3 Click **Add** > **Security Zone** or **Add** > **Interface Group**.

Step 4 Enter a **Name**.

Do not use the same name as a network or port object. These object names are deployed to the device, and duplicate names result in a failed deployment.

Step 5 Choose an **Interface Type**.

Step 6 (Optional) From the **Device** > **Interfaces** drop-down list, choose a device that contains interfaces you want to add.

You do not need to assign interfaces on this screen; you can instead assign interfaces to the zone or group when you configure the interface.

Step 7 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Enable the Physical Interface and Configure Ethernet Settings

This section describes how to:

- Enable the physical interface. By default, physical interfaces are disabled (with the exception of the Management interface).
- Set a specific speed and duplex. By default, speed and duplex are set to Auto.

This procedure only covers a small subset of Interface settings. Refrain from setting other parameters at this point. For example, you cannot name an interface that you want to use as part of an EtherChannel interface.

Note For the Firepower 4100/9300, you configure basic interface settings in FXOS. See Configure a Physical Interface, on page 332 for more information.

Note For Firepower 1010 and Secure Firewall 1210/1220 switch ports, see Configure Firepower 1010 and Secure Firewall 1210/1220 Switch Ports, on page 818.

You can also click the **Virtual Tunnels** tab to view details of dynamic and static VTIs of route-based VPNs on the device. For more information, see View Virtual Tunnel Information, on page 1555.

Before you begin

If you changed the physical interfaces on the device after you added it to the Firewall Management Center, you need to refresh the interface listing by clicking **Sync Interfaces from device** on the top left of **Interfaces**. For the Secure Firewall 3100/4200, which supports hot swapping, see Manage the Network Module for the Secure Firewall 3100/4200, on page 790 before you change interfaces on a device.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Enable the interface by checking the **Enabled** check box.

Step 4 (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

Step 5 (Optional) Set the duplex and speed by clicking **Hardware Configuration** > **Speed**.

- **Duplex**—Choose **Full** or **Half**. SFP interfaces only support **Full** duplex.
- **Speed**—Choose a speed (varies depending on the model). (Secure Firewall 3100/4200 only) Choose **Detect SFP** to detect the speed of the installed SFP module and use the appropriate speed. Duplex is always Full, and auto-negotiation is always enabled. This option is useful if you later change the network module to a different model, and want the speed to update automatically.For Secure Firewall 1250, you can configure a maximum interface speed of 2.5gbps.

  Note
  You cannot modify the speed of a HA or a Cluster Control Link interface.

- **Auto-negotiation**—Set the interface to negotiate the speed, link status, and flow control.
- **Forward Error Correction Mode**—(Secure Firewall 3100/4200 only) For 25 Gbps and higher interfaces, enable Forward Error Correction (FEC). For an EtherChannel member interface, you must configure FEC before you add it to the EtherChannel. The setting chosen when you use **Auto** depends on the transceiver type and whether the interface is fixed (built-in) or on a network module.

Table 50: Default FEC for Auto Setting

Transceiver Type	Fixed Port Default FEC (Ethernet 1/9 through 1/16)	Network Module Default FEC
25G-SR	Clause 108 RS-FEC	Clause 108 RS-FEC

Transceiver Type	Fixed Port Default FEC (Ethernet 1/9 through 1/16)	Network Module Default FEC
25G-LR	Clause 108 RS-FEC	Clause 108 RS-FEC
10/25G-CSR	Clause 108 RS-FEC	Clause 74 FC-FEC
25G-AOC*x*M	Clause 74 FC-FEC	Clause 74 FC-FEC
25G-CU2.5/3M	Auto-Negotiate	Auto-Negotiate
25G-CU4/5M	Auto-Negotiate	Auto-Negotiate
25/50/100G	Clause 91 RS-FEC	Clause 91 RS-FEC

Step 6 (Optional) (Firepower 1100/Secure Firewall 1200/3100/4200) Enable Link Layer Discovery Protocol (LLDP) by clicking **Hardware Configuration** > **Network Connectivity**.

- **Enable LLDP Receive**—Enables the firewall to receive LLDP packets from its peers.
- **Enable LLDP Transmit**—Enables the firewall to send LLDP packets to its peers.

Step 7 (Optional) (Secure Firewall 3100/4200) Enable pause (XOFF) frames for flow control by clicking **Hardware Configuration** > **Network Connectivity**, and checking **Flow Control Send**.

Flow control enables connected Ethernet ports to control traffic rates during congestion by allowing congested nodes to pause link operation at the other end. If the threat defense port experiences congestion (exhaustion of queuing resources on the internal switch) and cannot receive any more traffic, it notifies the other port by sending a pause frame to stop sending until the condition clears. Upon receipt of a pause frame, the sending device stops sending any data packets, which prevents any loss of data packets during the congestion period.

Note
The Firewall Threat Defense supports transmitting pause frames so that the remote peer can rate-control the traffic.

However, receiving of pause frames is not supported.

The internal switch has a global pool of 8000 buffers of 250 bytes each, and the switch allocates buffers dynamically to each port. A pause frame is sent out every interface with flowcontrol enabled when the buffer usage exceeds the global high-water mark (2 MB (8000 buffers)); and a pause frame is sent out of a particular interface when its buffer exceeds the port high-water mark (.3125 MB (1250 buffers)). After a pause is sent, an XON frame can be sent when the buffer usage is reduced below the low-water mark (1.25 MB globally (5000 buffers); .25 MB per port (1000 buffers)). The link partner can resume traffic after receiving an XON frame.

Only flow control frames defined in 802.3x are supported. Priority-based flow control is not supported.

Step 8 In the **Mode** drop-down list, choose one of the following:.

- **None**—Choose this setting for regular firewall interfaces and inline sets. The mode will automatically be changed to Routed, Switched, or Inline based on further configuration.
- **Passive**—Choose this setting for passive IPS-only interfaces.
- **Erspan**—Choose this setting for ERSPAN passive IPS-only interfaces.

Step 9 In the **Priority** field, enter a number ranging from 0–65535.

This value is used in the policy based routing configuration. The priority is used to determine how you want to distribute the traffic across multiple egress interfaces.

Step 10 Click **OK**.

Step 11 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Step 12 Continue configuring interfaces.

- Regular Firewall Interfaces, on page 817
- Inline Sets and Passive Interfaces, on page 893

Configure EtherChannel Interfaces

This section tells how to configure EtherChannel interfaces.

Note For the Firepower 4100/9300, you configure EtherChannels in FXOS. See Add an EtherChannel (Port Channel), on page 333 for more information.

About EtherChannels

This section describes EtherChannels.

About EtherChannels

An 802.3ad EtherChannel is a logical interface (called a port-channel interface) consisting of a bundle of individual Ethernet links (a channel group) so that you increase the bandwidth for a single network. A port channel interface is used in the same way as a physical interface when you configure interface-related features.

You can configure up to 48 EtherChannels, depending on how many interfaces your model supports.

Channel Group Interfaces

Each channel group can have up to 8 active interfaces, except for the ISA 3000, which supports 16 active interfaces. For switches that support only 8 active interfaces, you can assign up to 16 interfaces to a channel group: while only 8 interfaces can be active, the remaining interfaces can act as standby links in case of interface failure.

All interfaces in the channel group must be the same type and speed. The first interface added to the channel group determines the correct type and speed.

The EtherChannel aggregates the traffic across all the available active interfaces in the channel. The interface is selected using a proprietary hash algorithm, based on source or destination MAC addresses, IP addresses, TCP and UDP port numbers and VLAN numbers.

Connecting to an EtherChannel on Another Device

The device to which you connect the Firewall Threat Defense EtherChannel must also support 802.3ad EtherChannels; for example, you can connect to the Catalyst 6500 switch or the Cisco Nexus 7000.

When the switch is part of a Virtual Switching System (VSS) or Virtual Port Channel (vPC), then you can connect Firewall Threat Defense interfaces within the same EtherChannel to separate switches in the VSS/vPC. The switch interfaces are members of the same EtherChannel port-channel interface, because the separate switches act like a single switch.

Figure 318: Connecting to a VSS/vPC

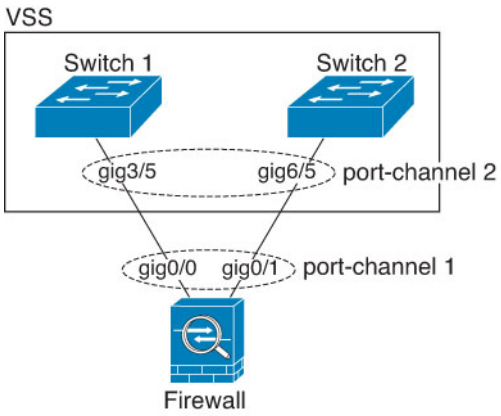

Note If the Firewall Threat Defense device is in transparent firewall mode, and you place the Firewall Threat Defense device between two sets of VSS/vPC switches, then be sure to disable Unidirectional Link Detection (UDLD) on any switch ports connected to the Firewall Threat Defense device with an EtherChannel. If you enable UDLD, then a switch port may receive UDLD packets sourced from both switches in the other VSS/vPC pair. The receiving switch will place the receiving interface in a down state with the reason "UDLD Neighbor mismatch".

If you use the Firewall Threat Defense device in an Active/Standby failover deployment, then you need to create separate EtherChannels on the switches in the VSS/vPC, one for each Firewall Threat Defense device. On each Firewall Threat Defense device, a single EtherChannel connects to both switches. Even if you could group all switch interfaces into a single EtherChannel connecting to both Firewall Threat Defense devices (in this case, the EtherChannel will not be established because of the separate Firewall Threat Defense system IDs), a single EtherChannel would not be desirable because you do not want traffic sent to the standby Firewall Threat Defense device.

Figure 319: Active/Standby Failover and VSS/vPC

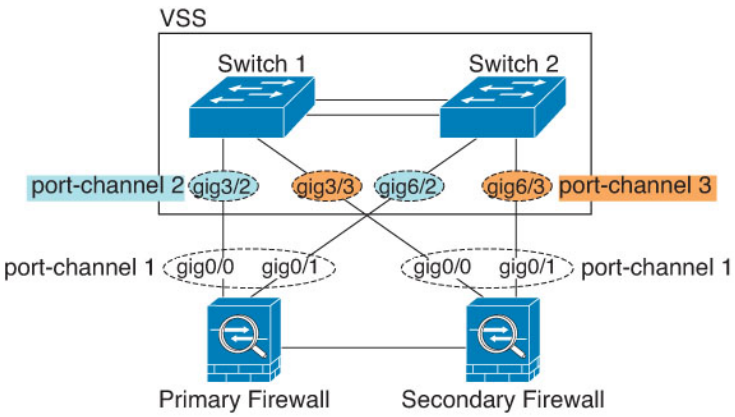

Link Aggregation Control Protocol

The Link Aggregation Control Protocol (LACP) aggregates interfaces by exchanging the Link Aggregation Control Protocol Data Units (LACPDUs) between two network devices.

You can configure each physical interface in an EtherChannel to be:

- Active—Sends and receives LACP updates. An active EtherChannel can establish connectivity with either an active or a passive EtherChannel. You should use the active mode unless you need to minimize the amount of LACP traffic.
- Passive—Receives LACP updates. A passive EtherChannel can only establish connectivity with an active EtherChannel. Not supported on hardware models.
- On—The EtherChannel is always on, and LACP is not used. An "on" EtherChannel can only establish a connection with another "on" EtherChannel.

LACP coordinates the automatic addition and deletion of links to the EtherChannel without user intervention. It also handles misconfigurations and checks that both ends of member interfaces are connected to the correct channel group. "On" mode cannot use standby interfaces in the channel group when an interface goes down, and the connectivity and configurations are not checked.

Load Balancing

The Firewall Threat Defense device distributes packets to the interfaces in the EtherChannel by hashing the source and destination IP address of the packet (this criteria is configurable). The resulting hash is divided by the number of active links in a modulo operation where the resulting remainder determines which interface owns the flow. All packets with a *hash_value* **mod** *active_links* result of 0 go to the first interface in the EtherChannel, packets with a result of 1 go to the second interface, packets with a result of 2 go to the third interface, and so on. For example, if you have 15 active links, then the modulo operation provides values from 0 to 14. For 6 active links, the values are 0 to 5, and so on.

If an active interface goes down and is not replaced by a standby interface, then traffic is rebalanced between the remaining links. The failure is masked from both Spanning Tree at Layer 2 and the routing table at Layer 3, so the switchover is transparent to other network devices.

EtherChannel MAC Address

All interfaces that are part of the channel group share the same MAC address. This feature makes the EtherChannel transparent to network applications and users, because they only see the one logical connection; they have no knowledge of the individual links.

Firepower and Secure Firewall Hardware

The port-channel interface uses the MAC address of the internal interface Internal-Data 0/1. Alternatively you can manually configure a MAC address for the port-channel interface. All EtherChannel interfaces on a chassis use the same MAC address, so be aware that if you use SNMP polling, for example, multiple interfaces will have the same MAC address.

Note Member interfaces only use the Internal-Data 0/1 MAC address after a reboot. Prior to rebooting, the member interface uses its own MAC address. If you add a new member interface after a reboot, you will have to perform another reboot to update its MAC address.

Guidelines for EtherChannels

Bridge Group

In routed mode, Firewall Management Center-defined EtherChannels are not supported as bridge group members. EtherChannels on the Firepower 4100/9300 can be bridge group members.

High Availability

- When you use an EtherChannel interface as a High Availability link, it must be pre-configured on both units in the High Availability pair; you cannot configure it on the primary unit and expect it to replicate to the secondary unit because *the High Availability link itself is required for replication*.
- If you use an EtherChannel interface for the state link, no special configuration is required; the configuration can replicate from the primary unit as normal. For the Firepower 4100/9300 chassis, all interfaces, including EtherChannels, need to be pre-configured on both units.
- You can monitor EtherChannel interfaces for High Availability. When an active member interface fails over to a standby interface, this activity does not cause the EtherChannel interface to appear to be failed when being monitored for device-level High Availability. Only when all physical interfaces fail does the EtherChannel interface appear to be failed (for an EtherChannel interface, the number of member interfaces allowed to fail is configurable).
- If you use an EtherChannel interface for a High Availability or state link, then to prevent out-of-order packets, only one interface in the EtherChannel is used. If that interface fails, then the next interface in the EtherChannel is used. You cannot alter the EtherChannel configuration while it is in use as a High Availability link. To alter the configuration, you need to temporarily disable High Availability, which prevents High Availability from occurring for the duration.

Model Support

- You cannot add EtherChannels in the Firewall Management Center for the Firepower 4100/9300 or the Firewall Threat Defense Virtual. The Firepower 4100/9300 supports EtherChannels, but you must perform all hardware configuration of EtherChannels in FXOS on the chassis.
- You cannot use Firepower 1010 or Secure Firewall 1210/1220 switch ports or VLAN interfaces in EtherChannels.

General EtherChannel Guidelines

- You can configure up to 48 EtherChannels, depending on how many interfaces are available on your model.
- Each channel group can have up to 8 active interfaces, except for the ISA 3000, which supports 16 active interfaces. For switches that support only 8 active interfaces, you can assign up to 16 interfaces to a channel group: while only 8 interfaces can be active, the remaining interfaces can act as standby links in case of interface failure.
- All interfaces in the channel group must be the same media type and speed capacity. The media type can be either RJ-45 or SFP; SFPs of different types (copper and fiber) can be mixed. You cannot mix interface capacities (for example 1GB and 10GB interfaces) by setting the speed to be lower on the larger-capacity interface, except for the Secure Firewall 1200/3100/4200, which supports different interface capacities as long as the speed is set to Detect SFP; in this case the lowest common speed is used.
- The device to which you connect the Firewall Threat Defense EtherChannel must also support 802.3ad EtherChannels.
- The Firewall Threat Defense device does not support LACPDUs that are VLAN-tagged. If you enable native VLAN tagging on the neighboring switch using the Cisco IOS **vlan dot1Q tag native** command, then the Firewall Threat Defense device will drop the tagged LACPDUs. Be sure to disable native VLAN tagging on the neighboring switch.
- The LACP rate depends on the model. When you set the rate (normal or fast), the device requests that rate from the connecting switch. In return, the device will send at the rate requested by the connecting switch. We recommend that you set the same rate on both sides.
    - Firepower 4100/9300—The LACP rate is set to fast by default in FXOS, but you can configure it as normal (also known as slow).
    - Secure Firewall 3100/4200—The LACP rate is set to normal (slow) by default, but you can configure it as fast on the device.
    - All other models—The LACP rate set to normal (also known as slow), and it is not configurable, which means the device will always request a slow rate from the connecting switch. We recommend setting the rate on the switch to slow, so both sides send LACP messages at the same rate.
- In Cisco IOS software versions earlier than 15.1(1)S2, Firewall Threat Defense did not support connecting an EtherChannel to a switch stack. With default switch settings, if the Firewall Threat Defense EtherChannel is connected cross stack, and if the primary switch is powered down, then the EtherChannel connected to the remaining switch will not come up. To improve compatibility, set the **stack-mac persistent timer** command to a large enough value to account for reload time; for example, 8 minutes or 0 for indefinite. Or, you can upgrade to more a more stable switch software version, such as 15.1(1)S2.

- All the Firewall Threat Defense configuration refers to the logical EtherChannel interface instead of the member physical interfaces.

Configure an EtherChannel

This section describes how to create an EtherChannel port-channel interface, assign interfaces to the EtherChannel, and customize the EtherChannel.

Guidelines

- You can configure up to 48 EtherChannels, depending on the number of interfaces for your model.
- Each channel group can have up to 8 active interfaces, except for the ISA 3000, which supports 16 active interfaces. For switches that support only 8 active interfaces, you can assign up to 16 interfaces to a channel group: while only 8 interfaces can be active, the remaining interfaces can act as standby links in case of interface failure.
- All interfaces in the channel group must be the same media type and speed capacity. The media type can be either RJ-45 or SFP; SFPs of different types (copper and fiber) can be mixed. You cannot mix interface capacities (for example 1GB and 10GB interfaces) by setting the speed to be lower on the larger-capacity interface, except for the Secure Firewall 1200/3100/4200, which supports different interface capacities as long as the speed is set to Detect SFP; in this case the lowest common speed is used.

Note For the Firepower 4100/9300, you configure EtherChannels in FXOS. See Add an EtherChannel (Port Channel), on page 333 for more information.

Before you begin

- You cannot add a physical interface to the channel group if you configured a name for it. You must first remove the name.

Note If you are using a physical interface already in your configuration, removing the name will clear any configuration that refers to the interface.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Enable the member interfaces according to Enable the Physical Interface and Configure Ethernet Settings, on page 778.

Step 3 Click **Add Interfaces** > **Ether Channel Interface**.

Step 4 On the **General** tab, set the **Ether Channel ID** to a number between 1 and 48 (1 and 8 for the Firepower 1010 and Secure Firewall 1210, 1 and 10 for the Secure Firewall 1220).

Figure 320: Add EtherChannel Interface

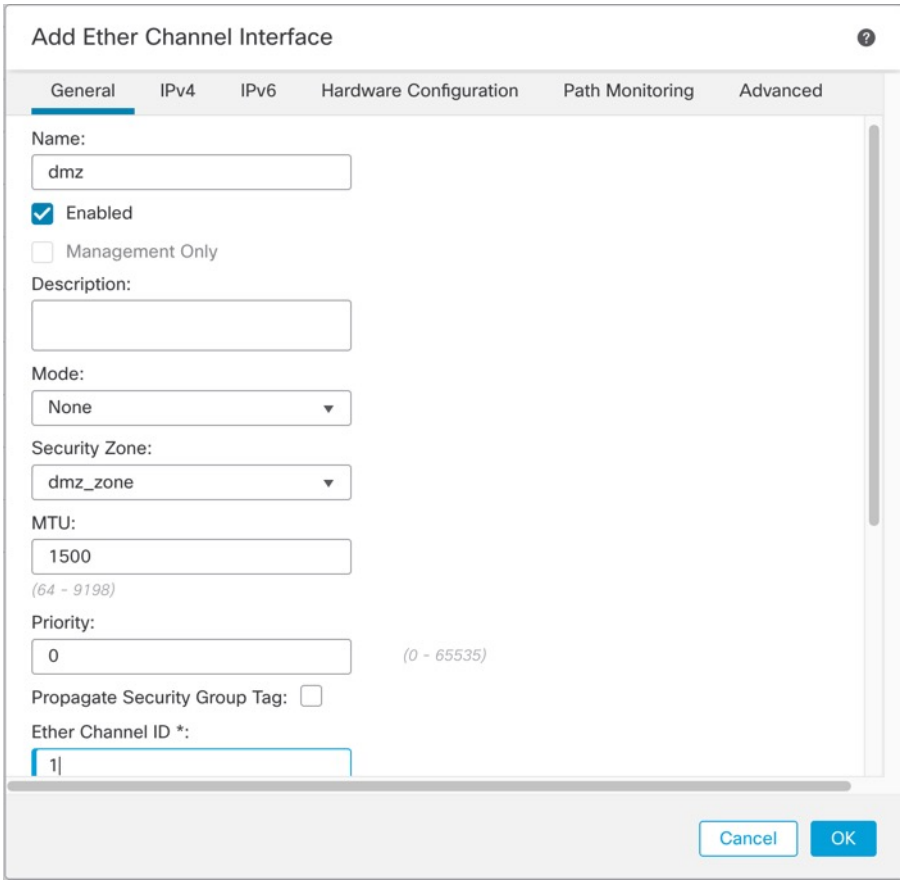

Step 5 In the **Available Interfaces** area, click an interface and then click **Add** to move it to the **Selected Interfaces** area. Repeat for all interfaces that you want to make members.

Make sure all interfaces are the same type and speed capability.

Figure 321: Available Interfaces

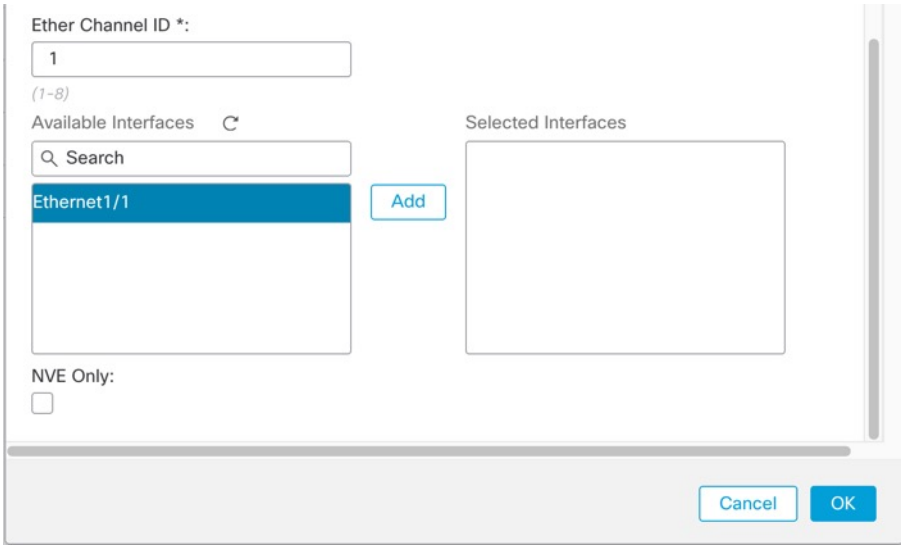

Step 6 (Optional) Click the **Advanced** tab to customize the EtherChannel. Set the following parameters on the **Information** sub-tab:

Figure 322: Advanced

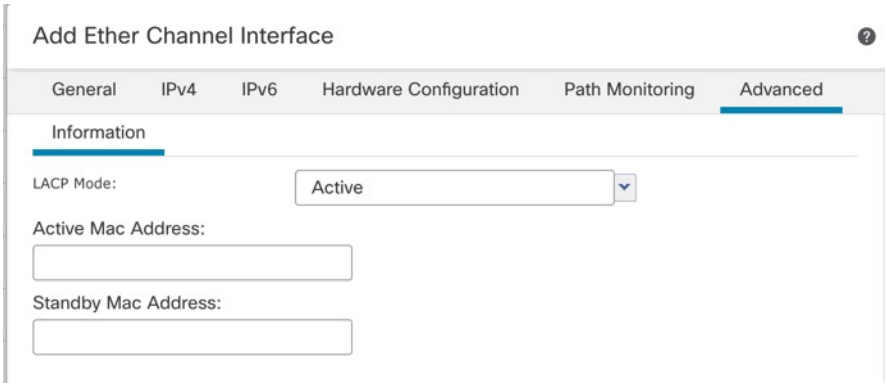

- (ISA 3000 only) **Load Balancing**—Select the criteria used to load balance the packets across the group channel interfaces. By default, the Firewall Threat Defense device balances the packet load on interfaces according to the source and destination IP address of the packet. If you want to change the properties on which the packet is categorized, choose a different set of criteria. For example, if your traffic is biased heavily towards the same source and destination IP addresses, then the traffic assignment to interfaces in the EtherChannel will be unbalanced. Changing to a different algorithm can result in more evenly distributed traffic. For more information about load balancing, see Load Balancing, on page 783.
- **LACP Mode**—Choose Active, Passive, or On. We recommend using Active mode (the default). Passive mode is only available for the ISA 3000 only.
- (Secure Firewall 3100/4200 only) **LACP Rate**—Choose Default, Normal, or Fast. The defualt is Normal (also known as slow). Sets the LACP data unit receive rate for a physical interface in the channel group. We recommend that you set the same rate on both sides.

- (ISA 3000 only) **Active Physical Interface: Range**—From the left drop-down list, choose the minimum number of active interfaces required for the EtherChannel to be active, between 1 and 16. The default is 1. From the right drop-down list, choose the maximum number of active interfaces allowed in the EtherChannel, between 1 and 16. The default is 16. If your switch does not support 16 active interfaces, be sure to set this command to 8 or fewer.
- **Active Mac Address**—Set a manual MAC address if desired. The mac_address is in H.H.H format, where H is a 16-bit hexadecimal digit. For example, the MAC address 00-0C-F1-42-4C-DE is entered as 000C.F142.4CDE.

Step 7 Click the **Hardware Configuration** tab and set the Duplex and Speed for all member interfaces.

For Secure Firewall 1250, you can configure a maximum interface speed of 2.5gbps.

Step 8 Click **OK**.

Step 9 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Step 10 (Optional) For regular firewall interfaces, add a VLAN subinterface. See Add a Subinterface, on page 836.

Step 11 For regular firewall interfaces, configure the routed or transparent mode interface parameters: Configure Routed Mode Interfaces, on page 856 or Configure Bridge Group Interfaces, on page 861. For IPS-only interfaces, see Inline Sets and Passive Interfaces, on page 893.

Sync Interface Changes with the Firewall Management Center

Addition or deletion of physical interfaces on the device can cause the Firewall Management Center and the device to get out of sync. The Firewall Management Center can detect interface changes by one of the following methods:

- Event sent from the device
- Sync when you deploy from the Firewall Management Center

  If the Firewall Management Center detects interface changes when it attempts to deploy, the deployment will fail. You must first accept the interface changes.
- Manual sync

Adding a new interface, or deleting an unused interface has minimal impact on the Firewall Threat Defense configuration. However, deleting an interface that is used in your security policy will impact the configuration. Interfaces can be referenced directly in many places in the Firewall Threat Defense configuration, including access rules, NAT, SSL, identity rules, VPN, DHCP server, and so on. Deleting an interface will delete any configuration associated with that interface. Policies that refer to security zones are not affected. You can also edit the membership of an allocated EtherChannel without affecting the logical device or requiring a sync on the Firewall Management Center.

When the Firewall Management Center detects changes, the **Interface** page shows status (removed, changed, or added) to the left of each interface.

This procedure describes how to manually sync interface changes if required. If interface changes are temporary, you should not save the changes in the Firewall Management Center; you should wait until the device is stable, and then re-sync.

Before you begin

- User Roles:
  - Admin
  - Access Admin
  - Network Admin

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 If required, click **Sync Interfaces** on the top left of **Interfaces**.

Step 3 After the changes are detected, see the following steps.

a) You will see a red banner on **Interfaces** indicating that the interface configuration has changed. Click the **Click to know more** link to view the interface changes.

b) Click **Validate Changes** to make sure your policy will still work with the interface changes.

If there are any errors, you need to change your policy and rerun the validation.

c) Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices.

Manage the Network Module for the Secure Firewall 3100/4200

If you install a network module before you first power on the device, no action is required; the network module is enabled and ready for use.

To view physical interface details for the device, and to manage the network module, open the **Chassis Operations** page. From **Devices** > **Device Management**, click **Manage** in the **Chassis** column. For clustering or High Availability, this option is only available for the control node/active unit. The **Chassis Operations** page opens for the device.

Figure 323: Chassis Operations

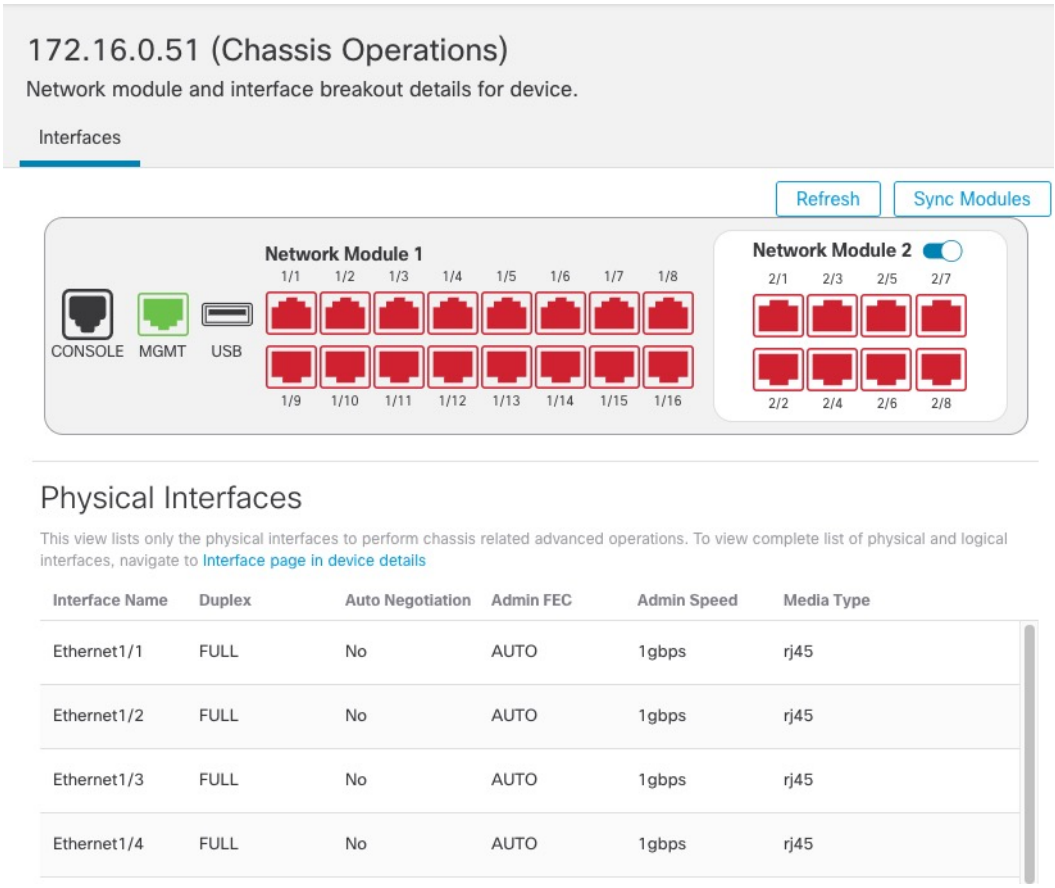

Click **Refresh** to refresh interface status. Click **Sync Modules** if you made a hardware change on the device that you need to detect.

If you need to make changes to your network module installation after initial bootup, then see the following procedures.

Configure Breakout Ports

You can configure 10GB breakout ports for each 40GB or higher interface. This procedure tells you how to break out and rejoin the ports. breakout ports can be used just like any other physical Ethernet port, including being added to EtherChannels.

Changes are immediate; you do not need to deploy to the device. After you break or rejoin, you cannot roll back to the previous interface state.

Before you begin

- You must use a supported breakout cable. See the hardware installation guide for more information.
- The interface cannot be in use for the following before breaking or rejoining:
  - Failover link

- Cluster control link
- Have a subinterface
- EtherChannel member
- BVI member
- Manager access interface

- Breaking or rejoining and interface that is used directly in your security policy can impact the configuration; however, the action is not blocked.

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column. For clustering or High Availability, this option is only available for the control node/active unit; network module changes are replicated to all nodes.

Figure 324: Manage Chassis

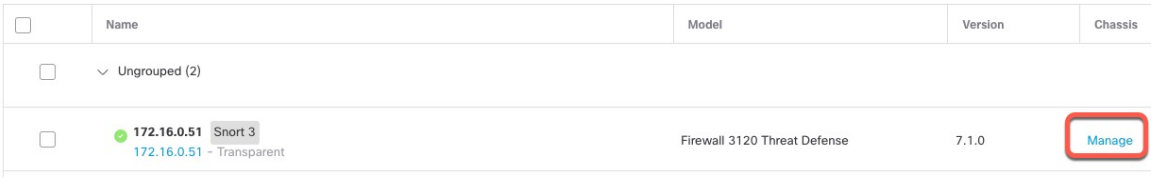

The **Chassis Operations** page opens for the device (in multi-instance mode, this page is called **Chassis Manager**). This page shows physical interface details for the device.

Step 2 Break out 10GB ports from a 40GB or higher interface.

a) click **Break** (↘) to the right of the interface.

Click **Yes** on the confirmation dialog box. If the interface is in use, you will see an error message. You must resolve any use cases before you can retry the breakout.

For example, to break out the Ethernet2/1 40GB interface, the resulting child interfaces will be identified as Ethernet2/1/1, Ethernet2/1/2, Ethernet2/1/3, and Ethernet2/1/4.

On the interfaces graphic, a port that is broken out has this appearance:

Figure 325: Breakout Ports

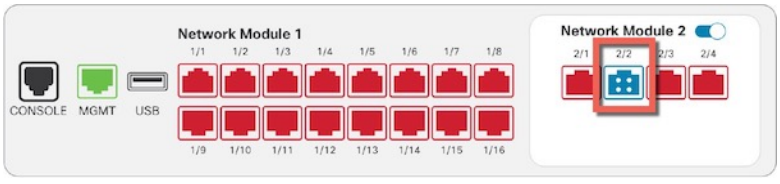

b) Click the link in the message at the top of the screen to go to the **Interfaces** page to save the interface changes.

Figure 326: Go to Interface Page

c) At the top of the **Interfaces** page, click **Click to know more**. The **Interface Changes** dialog box opens.

Figure 327: View Interface Changes

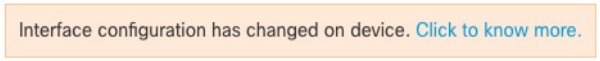

Figure 328: Interface Changes

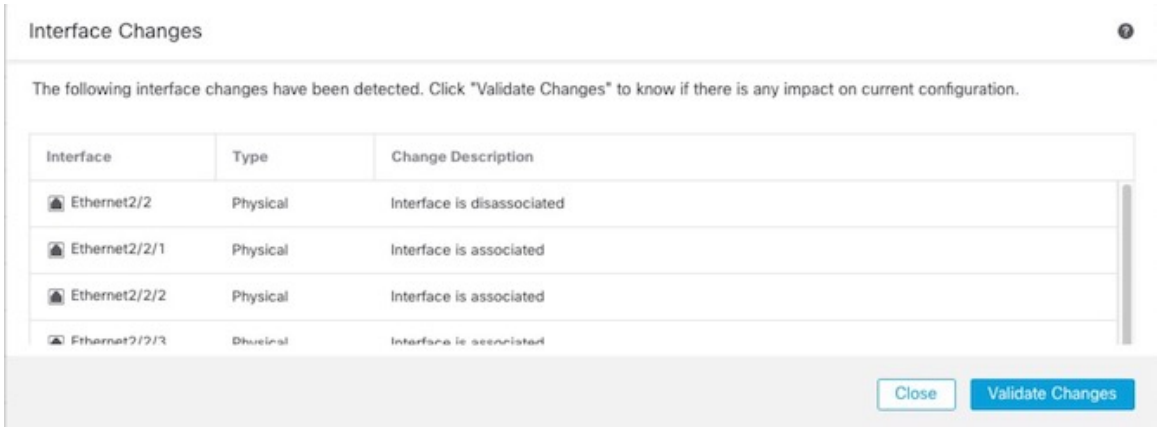

d) Click **Validate Changes** to make sure your policy will still work with the interface changes.

If there are any errors, you need to change your policy and rerun the validation.

Replacing the parent interface that is used in your security policy can impact the configuration. Interfaces can be referenced directly in many places in the configuration, including access rules, NAT, SSL, identity rules, VPN, DHCP server, and so on. Deleting an interface will delete any configuration associated with that interface. Policies that refer to security zones are not affected.

e) Click **Close** to return to the **Interfaces** page.
f) Click **Save** to save the interface changes to the firewall.
g) If you had to change any configuration, go to **Deploy** > **Deployment** and deploy the policy.

You do not need to deploy just to save the breakout port changes.

Step 3 Rejoin breakout ports.

You must rejoin all child ports for the interface.

a) Click **Join** ( ) to the right of the interface.

Click **Yes** on the confirmation dialog box. If any child ports are in use, you will see an error message. You must resolve any use cases before you can retry the rejoin.

b) Click the link in the message at the top of the screen to go to the **Interfaces** page to save the interface changes.

Figure 329: Go to Interface Page

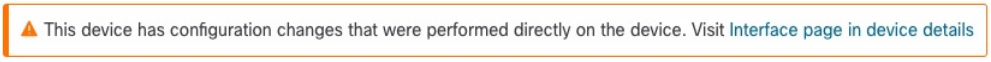

c) At the top of the **Interfaces** page, click **Click to know more**. The **Interface Changes** dialog box opens.

Figure 330: View Interface Changes

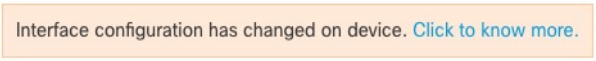

Figure 331: Interface Changes

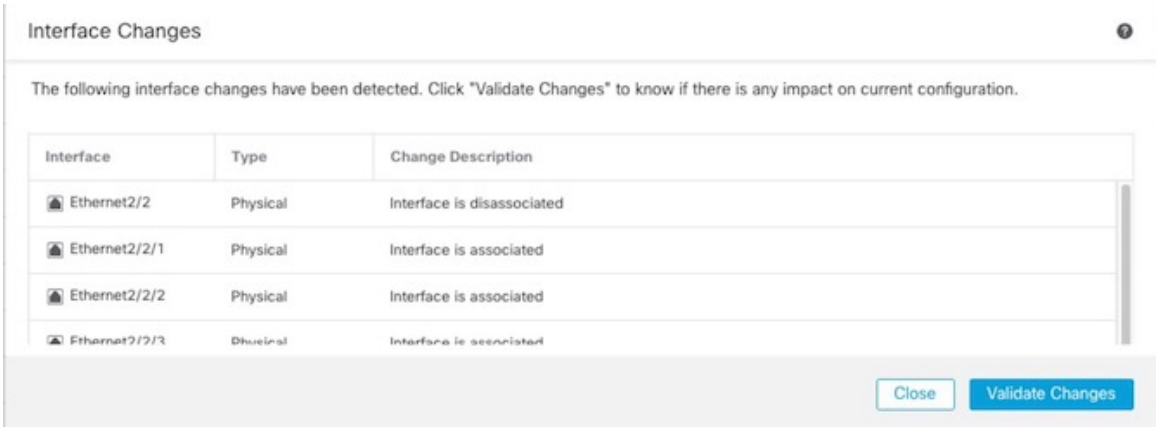

d) Click **Validate Changes** to make sure your policy will still work with the interface changes.

If there are any errors, you need to change your policy and rerun the validation.

Replacing the child interfaces that are used in your security policy can impact the configuration. Interfaces can be referenced directly in many places in the configuration, including access rules, NAT, SSL, identity rules, VPN, DHCP server, and so on. Deleting an interface will delete any configuration associated with that interface. Policies that refer to security zones are not affected.

e) Click **Close** to return to the **Interfaces** page.
f) Click **Save** to save the interface changes to the firewall.
g) If you had to change any configuration, go to **Deploy** > **Deployment** and deploy the policy.

You do not need to deploy just to save the breakout port changes.

Add a Network Module

To add a network module to a firewall after initial bootup, perform the following steps. Adding a new module requires a reboot.

Procedure

Step 1 Install the network module according to the hardware installation guide.

For clustering or High Availability, install the network module on all nodes.

Step 2 Reboot the firewall; see Shut Down or Restart the Device, on page 62.

For clustering or High Availability, reboot the data nodes/standby unit first, and wait for them to come back up. Then you can change the control node (see Change the Control Node, on page 494) or active unit (see Switch the Active Peer in the Firewall Threat Defense High Availability Pair, on page 446), and reboot the former control node/active unit.

Step 3 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column. For clustering or High Availability, this option is only available for the control node/active unit; network module changes are replicated to all nodes.

Figure 332: Manage Chassis

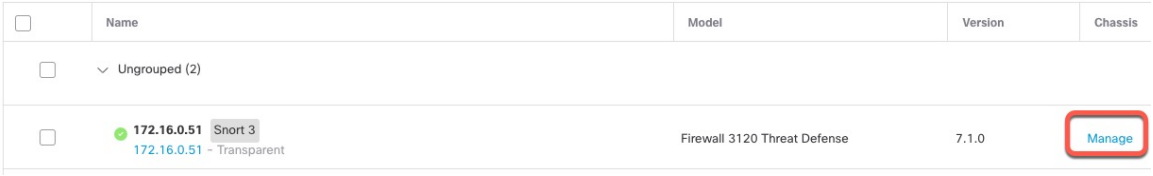

The **Chassis Operations** page opens for the device. This page shows physical interface details for the device.

Step 4 Click **Sync Modules** to update the page with the new network module details.

Step 5 On the interfaces graphic, click the slider ( ) to enable the network module.

Figure 333: Enable the Network Module

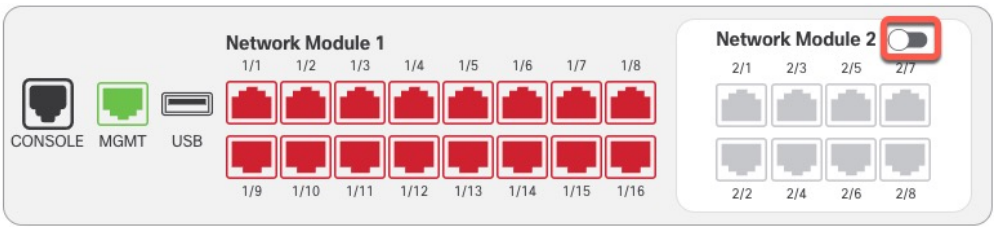

Step 6 You are prompted to confirm that you want to turn the network module on. Click **Yes**.

Figure 334: Confirm Enable

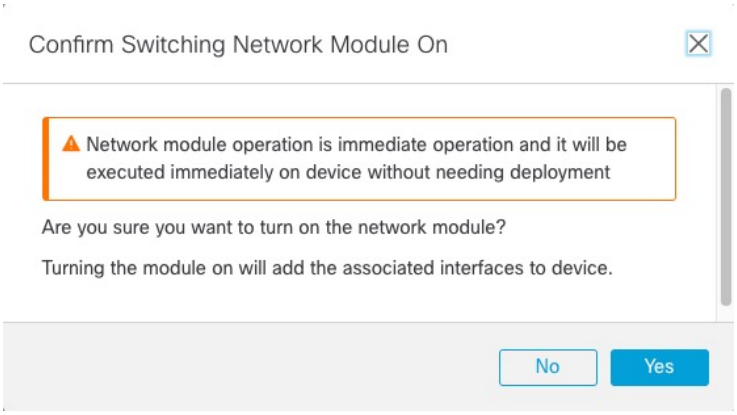

Step 7 You see a message at the top of the screen; click the link to go to the **Interfaces** page to save the interface changes.

Figure 335: Go to Interface Page

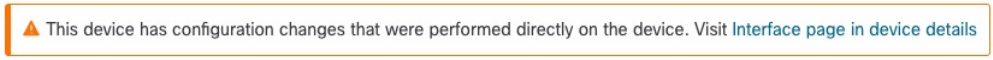

Step 8 (Optional) At the top of the **Interfaces** page, you see a message that the interface configuration has changed. You can click **Click to know more** to open the **Interface Changes** dialog box to view the changes.

Figure 336: View Interface Changes

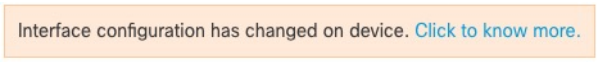

Figure 337: Interface Changes

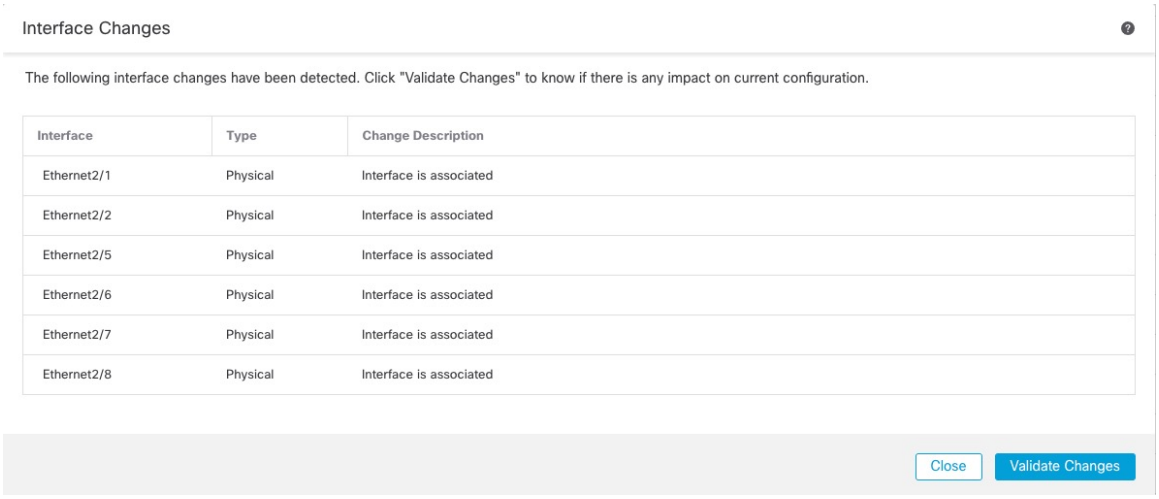

Click **Close** to return to the **Interfaces** page. (Because you are adding a new module, there shouldn't be any configuration impact, so you do not need to click **Validate Changes**.)

Step 9 Click **Save** to save the interface changes to the firewall.

Hot Swap the Network Module

You can hot swap a network module for a new module of the same type without having to reboot. However, you must shut down the current module to remove it safely. This procedure describes how to shut down the old module, install a new module, and enable it.

For clustering or High Availability, you can only perform chassis operations on the control node/active unit. You cannot disable a network module if the cluster control link/failover link is on the module.

Before you begin

Procedure

Step 1 For clustering or High Availability, perform the following steps.

- **Clustering**—Ensure the unit you want to perform the hot swap on is a data node (see Change the Control Node, on page 494); then break the node so it is no longer in the cluster. See Break a Node, on page 491.

  You will add the node back to the cluster after you perform the hot swap. Alternatively, you can perform all operations on the control node, and the network module changes will sync to all data nodes. However, you will lose use of those interfaces on all nodes during the hot swap.

- **High Availability**—To avoid failing over when you disable the network module:
  - If the failover link is on the network module, you must break High Availability. See Break a High Availability Pair, on page 450. Disabling the network module with an active failover link is not allowed.
  - Disable interface monitoring for interfaces on the network module. See Configure Standby IP Addresses and Interface Monitoring, on page 444.

Step 2 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column. For clustering or High Availability, this option is only available for the control node/active unit; network module changes are replicated to all nodes.

Figure 338: Manage Chassis

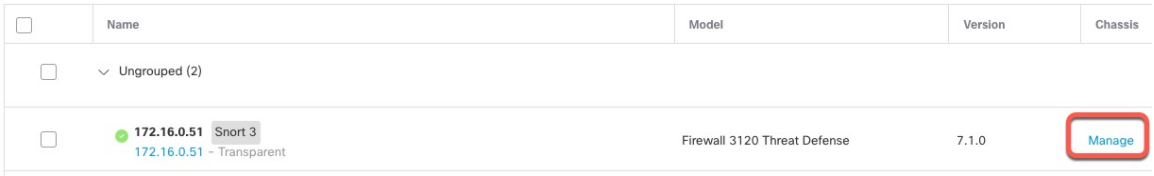

The **Chassis Operations** page opens for the device. This page shows physical interface details for the device.

Step 3 On the interfaces graphic, click the slider ( ) to disable the network module.

Figure 339: Disable the Network Module

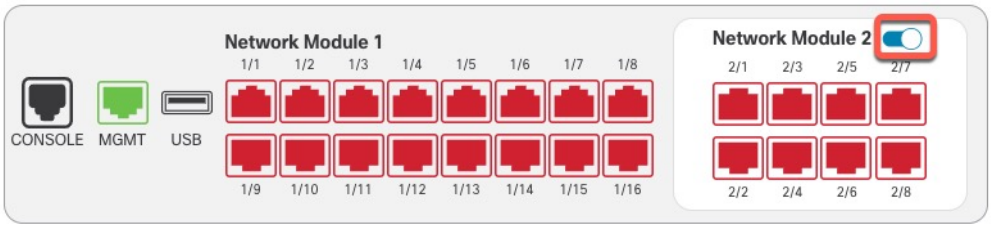

Do not save any changes on the **Interfaces** page. Because you are replacing the network module, you do not want to disrupt any existing configuration.

Step 4 You are prompted to confirm that you want to turn the network module off. Click **Yes**.

Figure 340: Confirm Disable

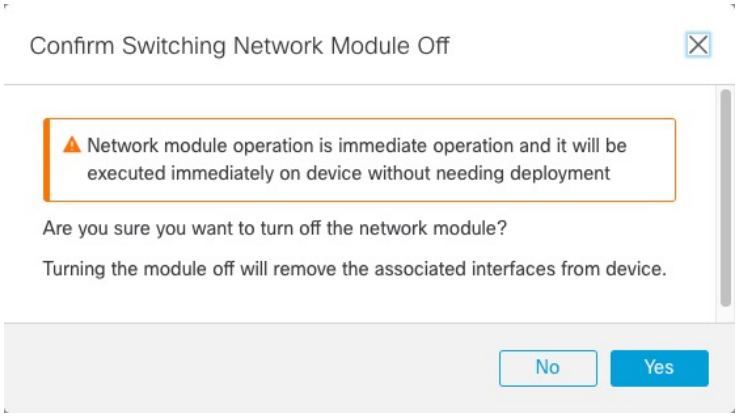

Step 5 On the device, remove the old network module and replace it with the new network module according to the hardware installation guide.

Step 6 In the Firewall Management Center, enable the new module by clicking the slider ( ).

Figure 341: Enable the Network Module

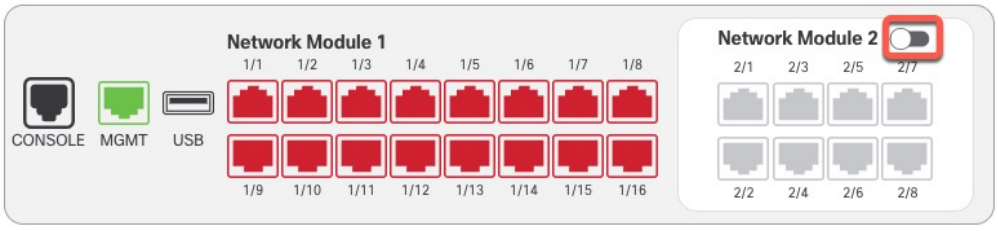

Step 7 You are prompted to confirm that you want to turn the network module on. Click **Yes**.

Figure 342: Confirm Enable

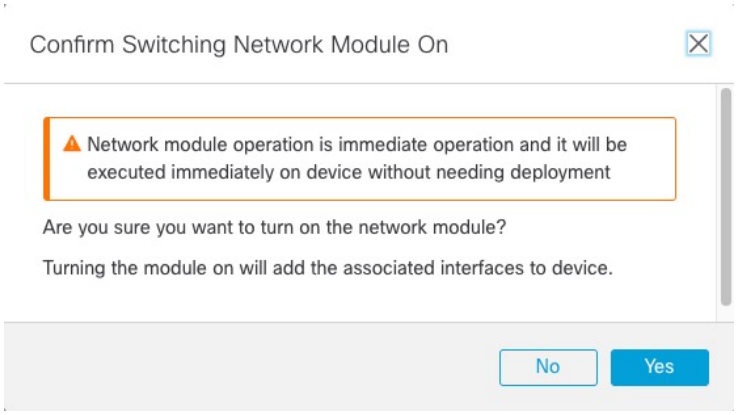

Step 8 For clustering or High Availability, perform the following steps.

- **Clustering**—Add the node back to the cluster. See Add a New Cluster Node, on page 489.
- **High Availability**—
  - If you broke High Availability, then reform High Availability. See Add a High Availability Pair, on page 441.
  - Reenable interface monitoring for interfaces on the network module. See Configure Standby IP Addresses and Interface Monitoring, on page 444.

Replace the Network Module with a Different Type

If you replace a network module with a different type, then a reboot is required. If the new module has fewer interfaces than the old module, you will have to manually remove any configuration related to interfaces that will no longer be present.

For clustering or High Availability, you can only perform chassis operations on the control node/active unit.

Before you begin

For High Availability, you cannot disable a network module if the failover link is on the module. You will have to break High Availability (see Break a High Availability Pair, on page 450), which means you will have downtime when you reboot the active unit. After the units finish rebooting, you can reform High Availability.

Procedure

Step 1 For clustering or High Availability, perform the following steps.

- **Clustering**—To avoid downtime, you can break each node one at a time so it is no longer in the cluster while you perform the network module replacement. See Break a Node, on page 491.

  You will add the node back to the cluster after you perform the replacement.

- **High Availability**—To avoid failing over when you replace the network module, disable interface monitoring for interfaces on the network module. See Configure Standby IP Addresses and Interface Monitoring, on page 444.

Step 2 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column. For clustering or High Availability, this option is only available for the control node/active unit; network module changes are replicated to all nodes.

Figure 343: Manage Chassis

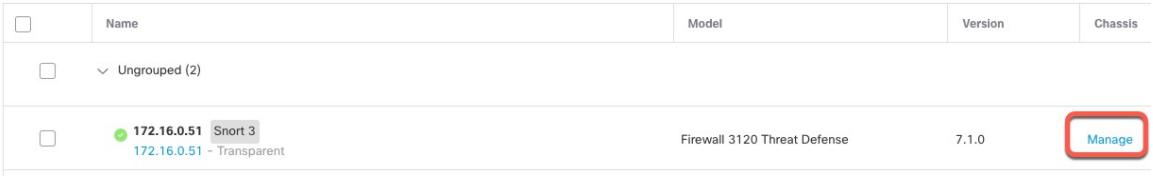

The **Chassis Operations** page opens for the device. This page shows physical interface details for the device.

Step 3 On the interfaces graphic, click the slider ( ) to disable the network module.

Figure 344: Disable the Network Module

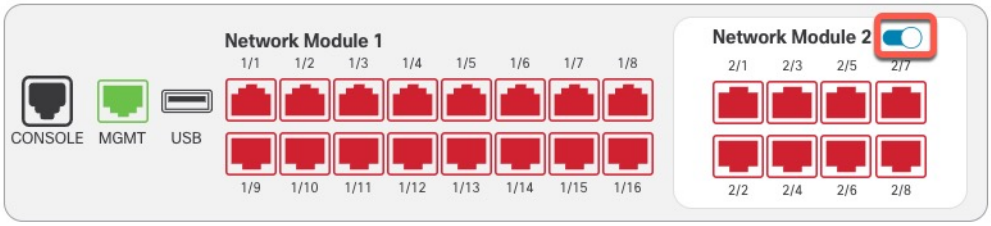

Do not save any changes on the **Interfaces** page. Because you are replacing the network module, you do not want to disrupt any existing configuration.

Step 4 You are prompted to confirm that you want to turn the network module off. Click **Yes**.

Figure 345: Confirm Disable

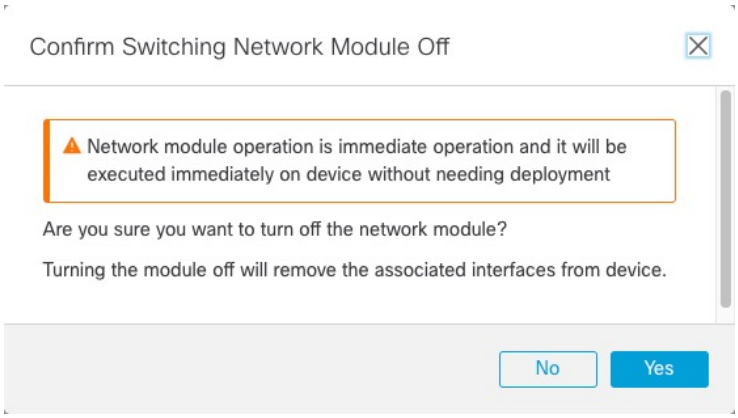

Step 5 On the device, remove the old network module and replace it with the new network module according to the hardware installation guide.

Step 6 Reboot the firewall; see Shut Down or Restart the Device, on page 62.

For clustering or High Availability, reboot the data nodes/standby unit first, and wait for them to come back up. Then you can change the control node (see Change the Control Node, on page 494) or active unit (see Switch the Active Peer in the Firewall Threat Defense High Availability Pair, on page 446), and reboot the former control node/active unit.

Step 7 In the Firewall Management Center, click **Sync Modules** to update the page with the new network module details.

Step 8 Enable the new module by clicking the slider ( ).

Figure 346: Enable the Network Module

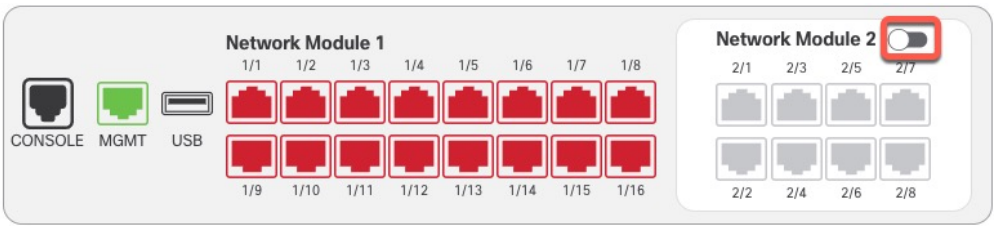

Step 9 You are prompted to confirm that you want to turn the network module on. Click **Yes**.

Figure 347: Confirm Enable

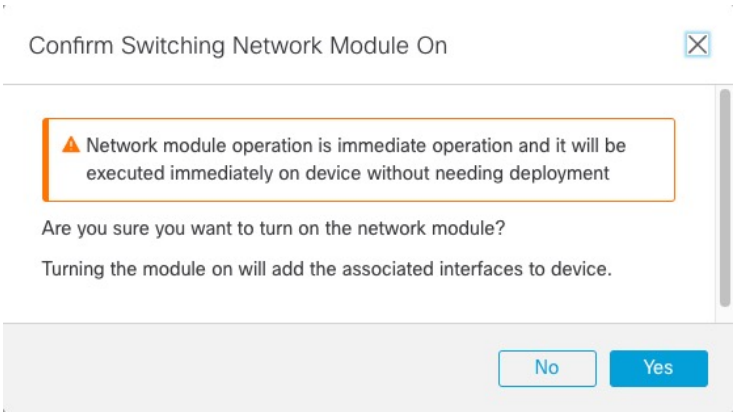

Step 10 Click the link in the message at the top of the screen to go to the **Interfaces** page to save the interface changes.

Figure 348: Go to Interface Page

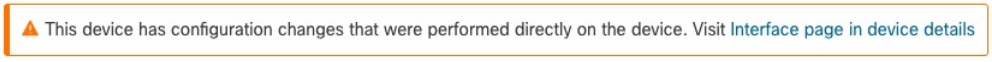

Step 11 If the network module has *fewer* interfaces:

a) At the top of the **Interfaces** page, click **Click to know more**. The **Interface Changes** dialog box opens.

Figure 349: View Interface Changes

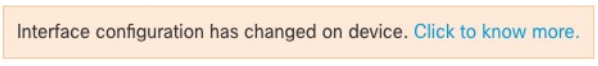

Figure 350: Interface Changes

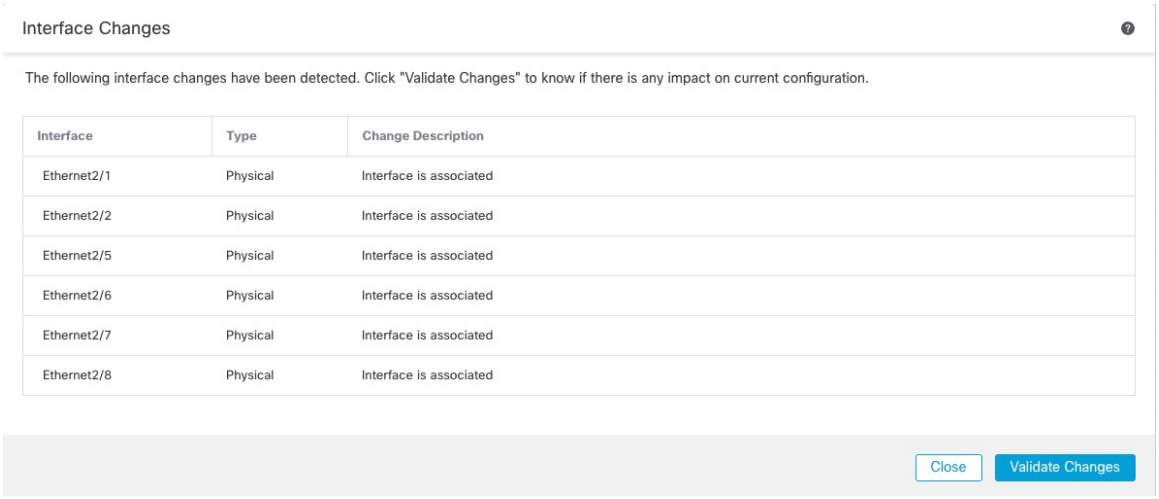

b) Click **Validate Changes** to make sure your policy will still work with the interface changes.

If there are any errors, you need to change your policy and rerun the validation.

Deleting an interface that is used in your security policy can impact the configuration. Interfaces can be referenced directly in many places in the configuration, including access rules, NAT, SSL, identity rules, VPN, DHCP server, and so on. Deleting an interface will delete any configuration associated with that interface. Policies that refer to security zones are not affected.

c) Click **Close** to return to the **Interfaces** page.

Step 12 To change the interface speed, see Enable the Physical Interface and Configure Ethernet Settings, on page 778.

The default speed is set to Detect SFP, which detects the correct speed from the SFP installed. You only need to fix the speed if you manually set the speed to a particular value and you now need a new speed.

Step 13 Click **Save** to save the interface changes to the firewall.

Step 14 If you had to change any configuration, go to **Deploy** > **Deployment** and deploy the policy.

You do not need to deploy just to save the network module changes.

Step 15 For clustering or High Availability, perform the following steps.

- **Clustering**—Add the node back to the cluster. See Add a New Cluster Node, on page 489.
- **High Availability**—Reenable interface monitoring for interfaces on the network module. See Configure Standby IP Addresses and Interface Monitoring, on page 444.

Remove the Network Module

If you want to permanently remove the network module, follow these steps. Removing a network module requires a reboot.

For clustering or High Availability, you can only perform chassis operations on the control node/active unit.

Before you begin

For clustering or High Availability, make sure the cluster/failover link is not on the network module.

Procedure

Step 1 From **Devices** > **Device Management**, click **Manage** in the **Chassis** column. For clustering or High Availability, this option is only available for the control node/active unit; network module changes are replicated to all nodes.

Figure 351: Manage Chassis

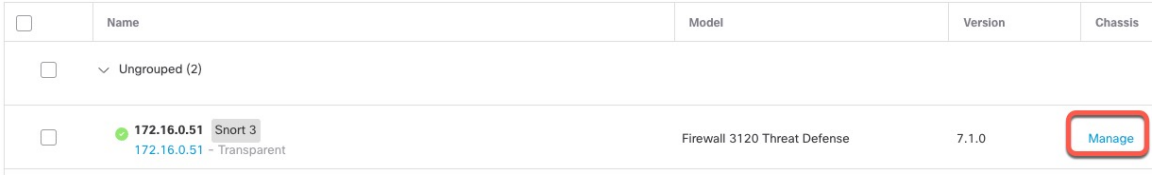

The **Chassis Operations** page opens for the device. This page shows physical interface details for the device.

Step 2 On the interfaces graphic, click the slider ( ) to disable the network module.

Figure 352: Disable the Network Module

Step 3 You are prompted to confirm that you want to turn the network module off. Click **Yes**.

Figure 353: Confirm Disable

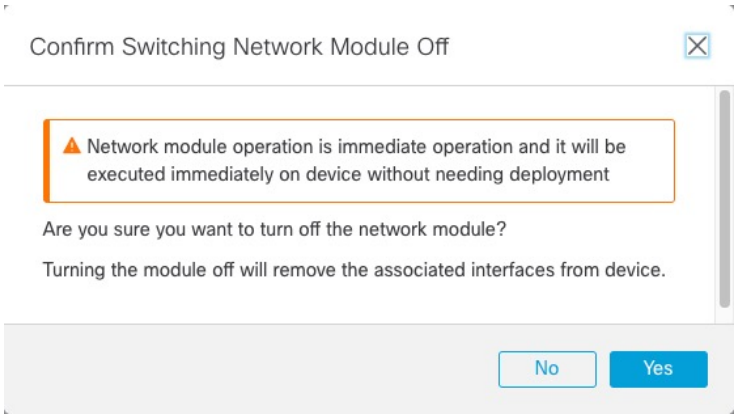

Step 4 You see a message at the top of the screen; click the link to go to the **Interfaces** page to save the interface changes.

Figure 354: Go to Interface Page

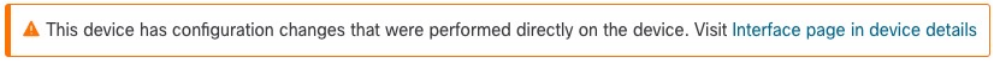

Step 5 At the top of the **Interfaces** page, you see a message that the interface configuration has changed.

Figure 355: View Interface Changes

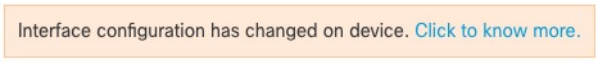

a) Click **Click to know more** to open the **Interface Changes** dialog box to view the changes.

Figure 356: Interface Changes

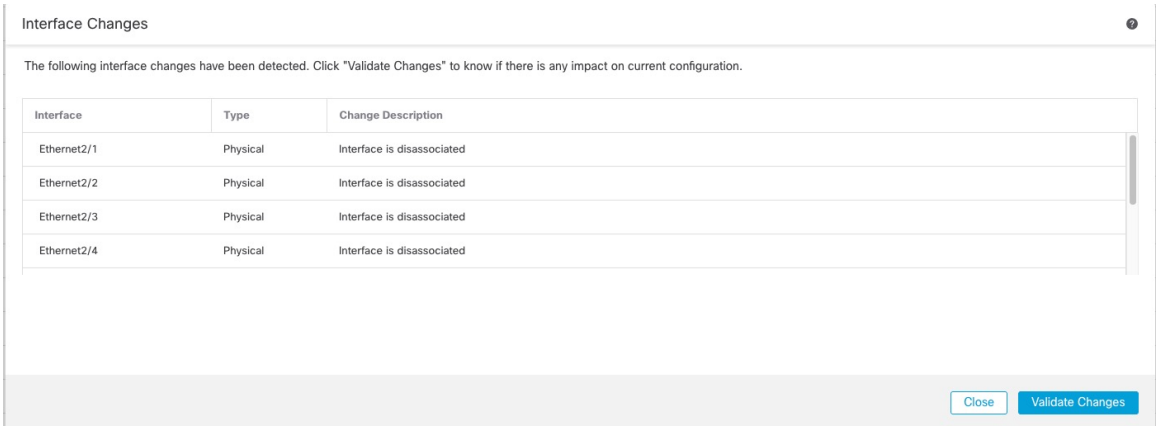

b) Click **Validate Changes** to make sure your policy will still work with the interface changes.

If there are any errors, you need to change your policy and rerun the validation.

Deleting an interface that is used in your security policy can impact the configuration. Interfaces can be referenced directly in many places in the configuration, including access rules, NAT, SSL, identity rules, VPN, DHCP server, and so on. Deleting an interface will delete any configuration associated with that interface. Policies that refer to security zones are not affected.

c) Click **Close** to return to the **Interfaces** page.

Step 6 Click **Save** to save the interface changes to the firewall.

Step 7 If you had to change any configuration, go to **Deploy** > **Deployment** and deploy the policy.

Step 8 Reboot the firewall; see Shut Down or Restart the Device, on page 62.

For clustering or High Availability, reboot the data nodes/standby unit first, and wait for them to come back up. Then you can change the control node (see Change the Control Node, on page 494) or active unit (see Switch the Active Peer in the Firewall Threat Defense High Availability Pair, on page 446), and reboot the former control node/active unit.

Merge the Management and Diagnostic Interfaces

Firewall Threat Defense 7.4 and later supports a merged Management and Diagnostic interface. If you have any configuration using the Diagnostic interface, then the interfaces will not be merged automatically, and you will need to perform the following procedure. This procedure requires you to acknowledge configuration changes, and in some cases, manually fix the configuration.

The Backup/Restore and Firewall Management Center configuration rollback functions save and restore the merged state, either non-merged or merged. For example, if you merge the interfaces, and then restore an old non-merged configuration, then the restored configuration will be in a non-merged state.

The following table shows the available configuration on the legacy Diagnostic interface, and how the merge is completed.

Table 51: Firewall Management Center Merged Management Interface Support

Legacy Diagnostic Interface Configuration	Merge Behavior	Supported on Management?
Interfaces		The "management" interface is now shown in read-only mode on the **Interfaces** page.
• IP address	Manual removal required.	The current Management IP address is used instead. For High Availability and clustering, the Management interface does not support a standby IP address or IP address pool; each unit has its own IP address that is maintained across failovers. Therefore, you cannot use a single management IP address to communicate with the current active/control unit. Set at the CLI using the **configure network ipv4** or **configure network ipv6** command.
• "diagnostic" name	Automatically changed to "management". **Note** No other interfaces can be named "management". You must change the name to proceed with the merge.	Changed to "management".

Legacy Diagnostic Interface Configuration	Merge Behavior	Supported on Management?
Static Routes	Manual removal required.	**No support.** The Management interface has a separate Linux routing table from the data interfaces. The Firewall Threat Defense actually has two "data" routing tables: for data interfaces and for management-only interfaces (which used to include Diagnostic, but also includes any interfaces you set to management-only). Depending on the traffic type, the Firewall Threat Defense checks one routing table, and then falls back to the other routing table. This route lookup no longer includes the Diagnostic interface, and does not include the Linux routing table for Management. See Routing Table for Management Traffic, on page 1163 for more information. You can add static routes for the Linux routing table at the CLI using the **configure network static-routes** command **Note** The *default* route is set with the **configure network ipv4** or **configure network ipv6** command.
Dynamic Routing	Manual removal required.	**No support.**
HTTP server	No change.	**No support.** This setting will no longer work on the merged device, but it is not removed from the Platform Settings. Platform Settings can be used for multiple devices, some of which may not yet be merged.
ICMP	No change.	**No support.** This setting will no longer work on the merged device, but it is not removed from the Platform Settings. Platform Settings can be used for multiple devices, some of which may not yet be merged.
Syslog Server	Automatically moved to Management interface.	Yes. The syslog server configuration already has the option to send syslogs out of the Management interface (starting in 6.3). If you had specifically chosen the Diagnostic interface for syslogs, it will be moved to use Management. **Note** If Platform Settings for syslog servers or SNMP hosts specify the Diagnostic interface by name, then you must use separate Platform Settings policies for merged and non-merged devices. **Note** The merged Management interface does not support Secure Syslogs.

Legacy Diagnostic Interface Configuration	Merge Behavior	Supported on Management?
SMTP	No change.	**No support.** The Firewall Threat Defense checks the data routing table only for the SMTP server, so you cannot use the Management interface or any other management-only interfaces. See Routing Table for Management Traffic, on page 1163 for more information.
SNMP	Automatically moved to Management interface.	Yes. The SNMP host configuration already has the option to allow SNMP hosts on the Management interface (starting in 6.3). If you had specifically chosen the Diagnostic interface for SNMP, it will be moved to use Management. **Note** If Platform Settings for syslog servers or SNMP hosts specify the Diagnostic interface by name, then you must use separate Platform Settings policies for merged and non-merged devices.
RADIUS server	Automatically moved to Management interface.	Yes. If you had specifically chosen the Diagnostic interface, it will be moved to use Management. **Note** If you specified a route lookup to find the source interface, then the Firewall Threat Defense will no longer be able to send traffic out of a management-only interface; you must explicitly select Management as the source interface. Other management-only interfaces cannot be used.
AD server	Automatically moved to Management interface.	Yes. If you had specifically chosen the Diagnostic interface, it will be moved to use Management. **Note** If you specified a route lookup to find the source interface, then the Firewall Threat Defense will no longer be able to send traffic out of a management-only interface; you must explicitly select Management as the source interface. Other management-only interfaces cannot be used.
DDNS	Manual removal required.	**No support.**
DHCP server	Manual removal required.	**No support.**

Legacy Diagnostic Interface Configuration	Merge Behavior	Supported on Management?
DNS server	Automatically moved to Management interface.	Yes. If you checked the **Enable DNS Lookup via diagnostic interface also** check box, then it will be moved to use Management. There is a routing lookup change when you do not choose any interfaces or check the **Enable DNS Lookup via diagnostic/management interface also** check box: the Firewall Threat Defense uses the data routing table only, and does not fall back to using the management-only routing table. Therefore, you cannot use a management-only interface for DNS other than the Management interface. **Note** The Management interface also has a separate DNS lookup setting for its management traffic only. Set at the CLI using the **configure network dns** command.
FlexConfig	Manual removal required.	**No support.**

Before you begin

- To view the current mode of the device, enter the **show management-interface convergence** command at the Firewall Threat Defense CLI. The following output shows that the Management interfaces are merged:

```
> show management-interface convergence
management-interface convergence
>
```

  The following output shows that the Management interfaces are not merged:

```
> show management-interface convergence
no management-interface convergence
>
```

- For High Availability pairs and clusters, perform this task on the active/control unit. The merged configuration will be replicated automatically to the standby/data units.

Procedure

Step 1 Choose **Devices** > **Device Management**, and click **Edit** (✎) for your Firewall Threat Defense. The **Interfaces** page is selected by default. .

Step 2 Edit the Diagnostic interface, and remove the IP address.

You cannot complete the merge until after you have removed the Diagnostic IP address.

Step 3 Click **Management Interface Merge** in the **Management Interface action needed** area.

The **Management Interface Merge** dialog box shows all the occurrences of the Diagnostic interface in the configuration. For any occurrences that require you to manually remove or change the configuration, they will appear with a warning icon. Platform Settings that will no longer work on your device are marked with a caution icon and require your acknowledgement.

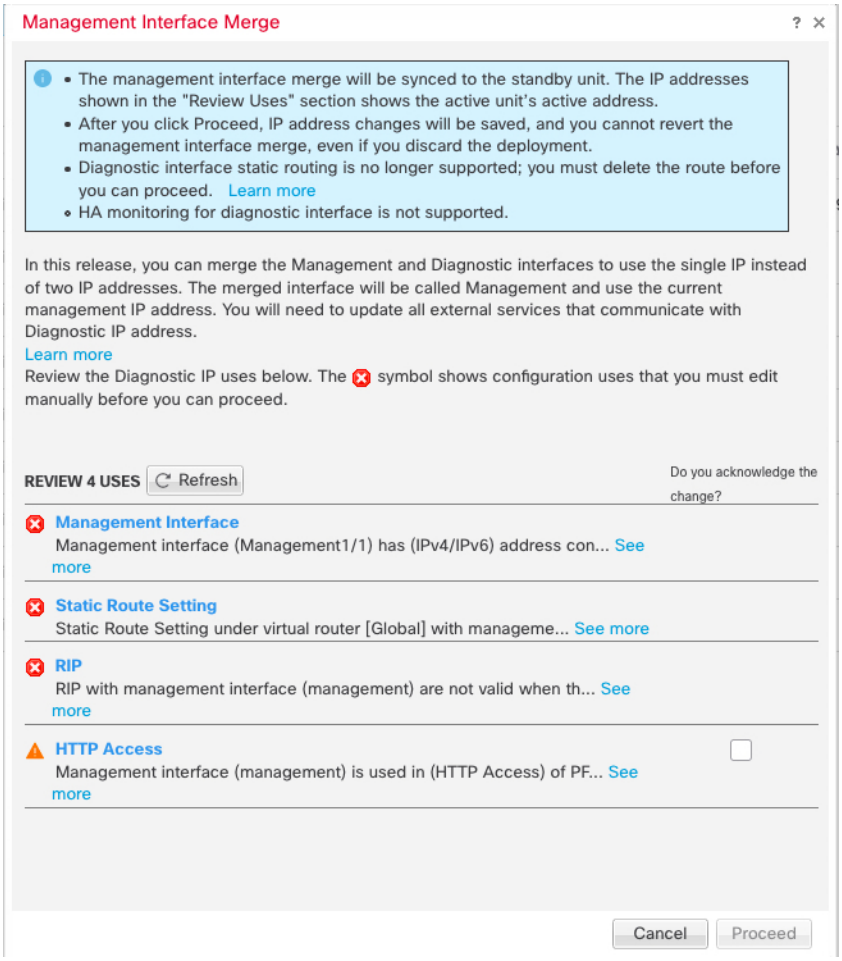

Step 4 If you need to manually remove or change any listed configurations, do the following.

a) Click **Cancel** to close the **Management Interface Merge** dialog box.

b) Navigate to the feature area. You can then delete the item, or choose a data interface instead.

c) Reopen the **Management Interface Merge** dialog box.

There should no longer be any warnings.

Step 5 For each configuration caution, click the box in **Do you acknowledge the change?** column, and then click **Proceed**.

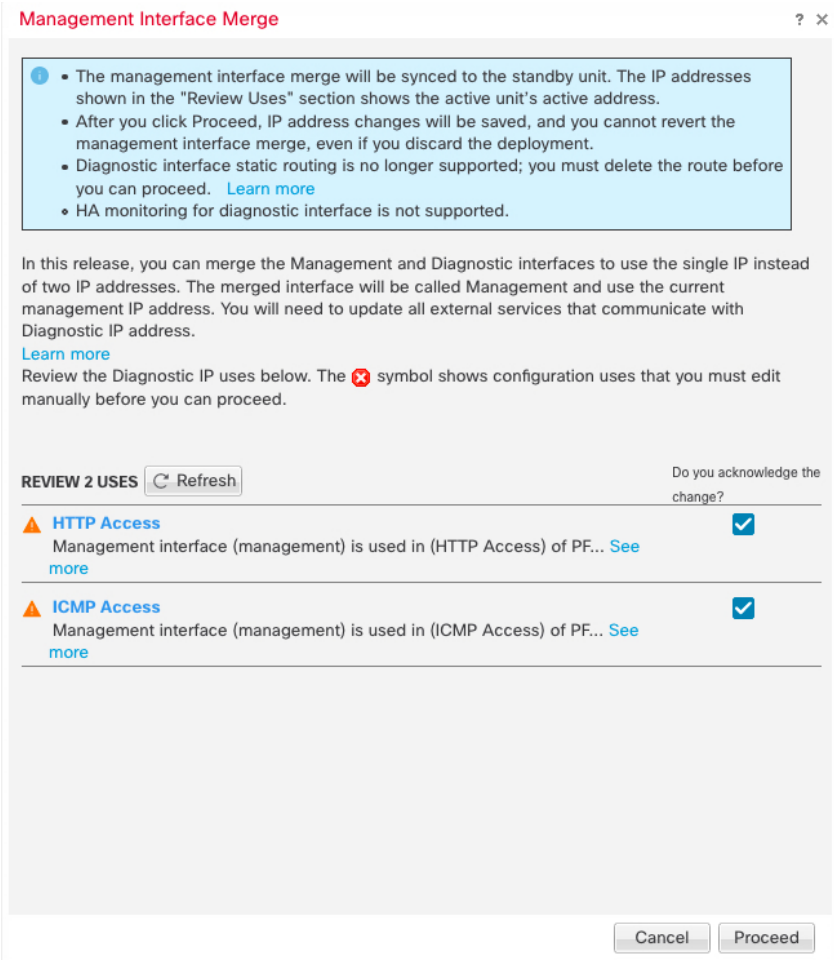

After the configuration is merged, you see a success banner:

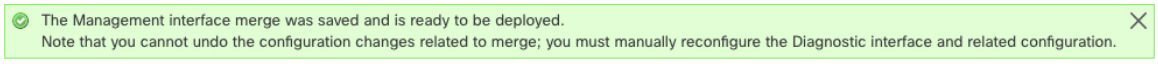

Step 6 Deploy the new merged configuration.

Caution
After you deploy the merged configuration, you can unmerge the interfaces from Firewall Management Center; however the Diagnostic interface will have to be reconfigured manually. See Unmerge the Management Interface, on page 811. Also, if you restore a configuration that is unmerged, or roll back to an unmerged configuration, then the device will revert to that unmerged configuration.

After the merge, the Management interface is shown on the **Interfaces** page, although it is read-only.

Step 7 After the merge, if you had any external services that communicated with the Diagnostic interface, you need to change their configuration to use the Management interface IP address.

For example:

- SNMP client

- RADIUS server—RADIUS servers often verify the IP address for incoming traffic, so you need to change that IP address to the Management address. Moreover, for a High Availability pair, you need to allow both the primary and secondary Management IP addresses; the Diagnostic interface used to support a single "floating" IP address that stayed with the active unit, but Management does not support that functionality.

Unmerge the Management Interface

The Firewall Threat Defense 7.4 and later supports a merged Management and Diagnostic interface. If you need to unmerge your interfaces, perform this procedure. We recommend using unmerged mode temporarily while you migrate your network to a merged mode deployment. Separate Management and Diagnostic interfaces may not be supported in all future releases.

Unmerging the interfaces does not restore your original Diagnostic configuration (if you upgraded and then merged your interfaces). You will need to reconfigure the Diagnostic interface manually. Also, the Management interface will now be named "management"; you cannot rename it "diagnostic."

Alternatively, if you used the Backup function to save an old unmerged configuration, you can restore that configuration or you can use the or Firewall Management Center configuration rollback feature, and the device will be in an unmerged state with the Diagnostic configuration intact.

Before you begin

- To view the current mode of the device, enter the **show management-interface convergence** command at the Firewall Threat Defense CLI. The following output shows that the Management interfaces are merged:

```
> show management-interface convergence
management-interface convergence
>
```

The following output shows that the Management interfaces are not merged:

```
> show management-interface convergence
no management-interface convergence
>
```

- For High Availability pairs and clusters, perform this task on the active/control unit. The merged configuration will be replicated automatically to the standby/data units.

Procedure

Step 1 Choose **Devices** > **Device Management**, and click **Edit** (✎) for your Firewall Threat Defense. The **Interfaces** page is selected by default. .

Step 2 For the Management interface, click **Unmerge Management Interface** (⇄).

Figure 357: Management Interface Selection

Step 3 Click **Yes** to confirm that you want to unmerge the interface.

Figure 358: Unmerge Confirmation

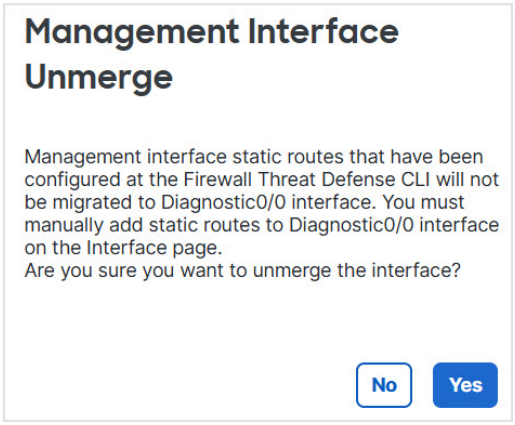

Step 4 Deploy the new unmerged configuration.

Note
If you restore a configuration that is merged, or roll back to a merged configuration, then the device will revert to that merged configuration.

After the merge, the Management interface is no longer shown on the **Interfaces** page.

History for Interfaces

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Sync Device is now called **Sync Interfaces**	7.7.0	7.7.0	**Sync Device** was changed to **Sync Interfaces** to indicate that this function is only for interface changes. This function no longer detects changes made to the manager access interface; see **Devices** > **Device Management** > **Device** > **Management** > **Manager Access Details: Configuration**. Other out-of-band configuration changes performed at the diagnostic CLI in recovery-config mode need to be discovered at **Devices** > **Device Management** > **Device** > **Health** > **Out of Band Status**. New/Modified screens: **Devices** > **Device Management** > **Interfaces**

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Loopback and Management type interface group objects	7.4.0	7.4.0	You can now create interface group objects that include only management-only interfaces or only loopback interfaces. You can then use these groups for management features such as DNS servers, HTTP access, or SSH. Loopback groups are supported for any feature that supports loopback interfaces. Note that DNS does not support management interfaces. New/Modified screens: **Objects** > **Object Management** > **Interface** > **Add** > **Interface Group**
Merged Management and Diagnostic interfaces	7.4.0	7.4.0	For new devices using 7.4 and later, you cannot use the legacy Diagnostic interface. Only the merged Management interface is available. If you upgraded to 7.4 or later, and you did not have any configuration for the Diagnostic interface, then the interfaces will merge automatically. If you upgraded to 7.4 or later, and you have configuration for the Diagnostic interface, then you have the choice to merge the interfaces manually, or you can continue to use the separate Diagnostic interface. Note that support for the Diagnostic interface will be removed in a later release, so you should plan to merge the interfaces as soon as possible. Merged mode also changes the behavior of AAA traffic to use the data routing table by default. The management-only routing table can now only be used if you specify the management-only interface (including Management) in the configuration. New/Modified screens: **Devices** > **Device Management** > **Interfaces** New/Modified commands: **show management-interface convergence**
Default Forward Error Correction (FEC) on Secure Firewall 3100 fixed ports changed to Clause 108 RS-FEC from Clause 74 FC-FEC for 25 GB+ SR, CSR, and LR transceivers	7.2.4	7.2.4	When you set the FEC to Auto on the Secure Firewall 3100 fixed ports, the default type is now set to Clause 108 RS-FEC instead of Clause 74 FC-FEC for 25 GB+ SR, CSR, and LR transceivers.
LLDP support for the Firepower 2100, Secure Firewall 3100	7.2.0	7.2.0	You can enable Link Layer Discovery Protocol (LLDP) for Firepower 2100 and Secure Firewall 3100 interfaces. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Hardware Configuration** > **Network Connectivity** New/Modified commands: **show lldp status, show lldp neighbors, show lldp statistics**

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Pause Frames for Flow Control for the Secure Firewall 3100	7.2.0	7.2.0	If you have a traffic burst, dropped packets can occur if the burst exceeds the buffering capacity of the FIFO buffer on the NIC and the receive ring buffers. Enabling pause frames for flow control can alleviate this issue. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Hardware Configuration** > **Network Connectivity**
Support for hot swapping the network module for the Secure Firewall 3100	7.1.0	7.1.0	You can add or remove the network module on the Secure Firewall 3100 while the firewall is powered up. To replace a module with another module of the same type, you do not need to reboot. After initial bootup, adding a module, permanently removing a module, or replacing a module with a new type requires a reboot. New/Modified screens: **Devices** > **Device Management** > **Chassis Operations**
Support for Forward Error Correction for the Secure Firewall 3100	7.1.0	7.1.0	Secure Firewall 3100 25 Gbps interfaces support Forward Error Correction (FEC). FEC is enabled by default and set to Auto. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** > **Hardware Configuration**
Support for setting the speed based on the SFP for the Secure Firewall 3100	7.1.0	7.1.0	The Secure Firewall 3100 supports speed detection for interfaces based on the SFP installed. Detect SFP is enabled by default. This option is useful if you later change the network module to a different model, and want the speed to update automatically. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** > **Hardware Configuration**
LLDP support for the Firepower 1100.	7.1.0	7.1.0	You can enable Link Layer Discovery Protocol (LLDP) for Firepower 1100 interfaces. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Hardware Configuration** > **LLDP** New/Modified commands: **show lldp status, show lldp neighbors, show lldp statistics**
Interface auto-negotiation is now set independently from speed and duplex, interface sync improved.	7.1.0	7.1.0	Interface auto-negotiation is now set independently from speed and duplex. Also, when you sync the interfaces in Firewall Management Center, hardware changes are detected more effectively. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Hardware Configuration** > **Speed** Supported platforms: Firepower 1000, 2100, Secure Firewall 3100

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Firepower 1100/2100 series fiber interfaces now support disabling auto-negotiation.	6.7.0	6.7.0	You can now configure a Firepower 1100/2100 series fiber interface to disable flow control and link status negotiation. Previously, when you set the fiber interface speed (1000 or 10000 Mbps) on these devices, flow control and link status negotiation was automatically enabled. You could not disable it. Now, you can deselect **Auto-negotiation** and set the speed to 1000 to disable flow control and link status negotiation. You cannot disable negotiation at 10000 Mbps. New/modified screens: **Devices** > **Device Management** > **Interfaces** > **Hardware Configuration** > **Speed**

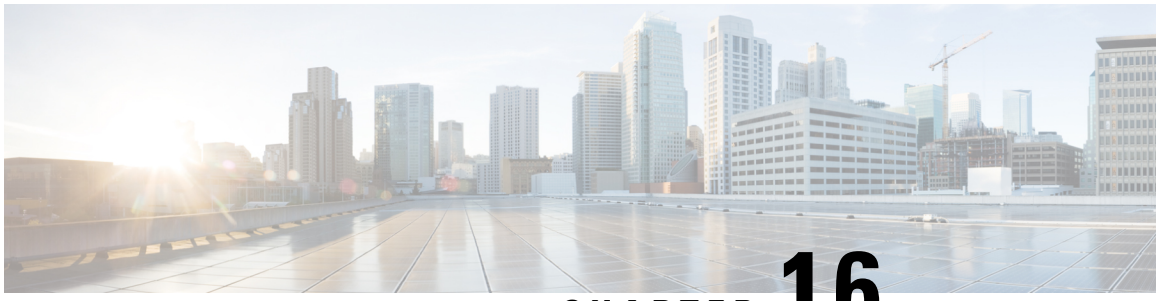

CHAPTER 16

Regular Firewall Interfaces

This chapter includes regular firewall Firewall Threat Defense interface configuration including EtherChannels, VLAN subinterfaces, IP addressing, and more.

Note For initial interface configuration on the Firepower 4100/9300, see Configure Interfaces, on page 332.

- Requirements and Prerequisites for Regular Firewall Interfaces, on page 817
- Configure Firepower 1010 and Secure Firewall 1210/1220 Switch Ports, on page 818
- Configure Loopback Interfaces, on page 828
- Configure VLAN Subinterfaces and 802.1Q Trunking, on page 834
- Configure VXLAN Interfaces, on page 838
- Configure Routed and Transparent Mode Interfaces, on page 853
- Configure Advanced Interface Settings, on page 875
- History for Regular Firewall Interfaces, on page 886

Requirements and Prerequisites for Regular Firewall Interfaces

Model Support

Threat Defense

User Roles

- Admin
- Access Admin
- Network Admin

Configure Firepower 1010 and Secure Firewall 1210/1220 Switch Ports

You can configure each Firepower 1010 and Secure Firewall 1210/1220 interface to run as a regular firewall interface or as a Layer 2 hardware switch port. This section includes tasks for starting your switch port configuration, including enabling or disabling the switch mode and creating VLAN interfaces and assigning them to switch ports. This section also describes how to customize Power over Ethernet (PoE) on supported interfaces.

About Switch Ports

This section describes the switch ports of Firepower 1010, and Secure Firewall 1210/1220.

Understanding Switch Ports and Interfaces

Ports and Interfaces

For each physical 1010/1210/1220 interface, you can set its operation as a firewall interface or as a switch port. See the following information about physical interface and port types as well as logical VLAN interfaces to which you assign switch ports:

- Physical firewall interface—In routed mode, these interfaces forward traffic between networks at Layer 3, using the configured security policy to apply firewall and VPN services. In transparent mode, these interfaces are bridge group members that forward traffic between the interfaces on the same network at Layer 2, using the configured security policy to apply firewall services. In routed mode, you can also use Integrated Routing and Bridging with some interfaces as bridge group members and others as Layer 3 interfaces. By default, the Ethernet 1/1 interface is configured as a firewall interface. You can also configure these interfaces to be IPS-only (inline sets and passive interfaces).
- Physical switch port—Switch ports forward traffic at Layer 2, using the switching function in hardware. Switch ports on the same VLAN can communicate with each other using hardware switching, and traffic is not subject to the Firewall Threat Defense security policy. Access ports accept only untagged traffic, and you can assign them to a single VLAN. Trunk ports accept untagged and tagged traffic, and can belong to more than one VLAN. By default, Ethernet 1/2 through 1/8 (1010 and 1210), or Ethernet 1/2 through 1/10 (1220) are configured as access switch ports on VLAN 1. You cannot configure the Management interface as a switch port.
- Logical VLAN interface—These interfaces operate the same as physical firewall interfaces, with the exception being that you cannot create subinterfaces, IPS-only interfaces (inline sets and passive interfaces), or EtherChannel interfaces. When a switch port needs to communicate with another network, then the Firewall Threat Defense device applies the security policy to the VLAN interface and routes to another logical VLAN interface or firewall interface. You can even use Integrated Routing and Bridging with VLAN interfaces as bridge group members. Traffic between switch ports on the same VLAN are not subject to the Firewall Threat Defense security policy, but traffic between VLANs in a bridge group are subject to the security policy, so you may choose to layer bridge groups and switch ports to enforce the security policy between certain segments.

Power Over Ethernet

PoE is available on the following ports:

- Firepower 1010—Ethernet 1/7 and 1/8 using IEEE 802.3af (PoE) and 802.3at (PoE+) up to 30 watts per port, up to a combined 60 watts.
- Secure Firewall 1210CP—Ethernet 1/5, 1/6, 1/7, and 1/8 using IEEE 802.3af (PoE), 802.3at (PoE+), and 802.3bt (PoE++ and Hi-PoE) up to 90 watts per port, up to a combined 120 watts.

Note PoE is not supported on the 1010E, 1210CE, and 1220CX.

PoE+ or higher uses Link Layer Discovery Protocol (LLDP) to negotiate the power level. Power is only supplied when needed.

If you shut down the interface, then you disable power to the device.

Auto-MDI/MDIX Feature

For all switch ports, the default auto-negotiation setting also includes the Auto-MDI/MDIX feature. Auto-MDI/MDIX eliminates the need for crossover cabling by performing an internal crossover when a straight cable is detected during the auto-negotiation phase. Either the speed or duplex must be set to auto-negotiate to enable Auto-MDI/MDIX for the interface. If you explicitly set both the speed and duplex to a fixed value, thus disabling auto-negotiation for both settings, then Auto-MDI/MDIX is also disabled. When the speed and duplex are set to 1000 and full, then the interface always auto-negotiates; therefore Auto-MDI/MDIX is always enabled and you cannot disable it.

Guidelines and Limitations for Switch Ports

High Availability and Clustering

- No cluster support.
- You should not use the switch port functionality when using High Availability. Because the switch ports operate in hardware, they continue to pass traffic on both the active *and* the standby units. High Availability is designed to prevent traffic from passing through the standby unit, but this feature does not extend to switch ports. In a normal High Availability network setup, active switch ports on both units will lead to network loops. We suggest that you use external switches for any switching capability. Note that VLAN interfaces can be monitored by failover, while switch ports cannot. Theoretically, you can put a single switch port on a VLAN and successfully use High Availability, but a simpler setup is to use physical firewall interfaces instead.
- You can only use a firewall interface as the failover link.

Logical VLAN Interfaces (SVIs)

- You can create up to 60 VLAN interfaces.
- If you also use VLAN subinterfaces on a firewall interface, you cannot use the same VLAN ID as for a logical VLAN interface. VLAN 1 is reserved for the logical VLAN interface.

- MAC Addresses:
  - Routed firewall mode—All VLAN interfaces share a MAC address. Ensure that any connected switches can support this scenario. If the connected switches require unique MAC addresses, you can manually assign MAC addresses. See Configure the MAC Address, on page 881.
  - Transparent firewall mode—Each VLAN interface has a unique MAC address. You can override the generated MAC addresses if desired by manually assigning MAC addresses. See Configure the MAC Address, on page 881.

Bridge Groups

You cannot mix logical VLAN interfaces and physical firewall interfaces in the same bridge group.

VLAN Interface and Switch Port Unsupported Features

VLAN interfaces and switch ports do not support:

- Dynamic routing
- Multicast routing
- Equal-Cost Multi-Path routing (ECMP)
- Inline sets or Passive interfaces
- EtherChannels—Switch ports cannot be part of an EtherChannel. PoE is also not supported on a port in an EtherChannel.
- Failover and state link
- Security group tagging (SGT)

Other Guidelines and Limitations

- You can configure a maximum of 60 named interfaces on the Firepower 1010 and Secure Firewall 1210/1220.
- You cannot configure the Management interface as a switch port.

Default Settings

- Ethernet 1/1 is a firewall interface.
- On 1010/1210, Ethernet 1/2 through Ethernet 1/8 are switch ports assigned to VLAN 1.
- On 1220, Ethernet 1/2 through Ethernet 1/10 are switch ports assigned to VLAN 1.
- Default Speed and Duplex—By default, the speed and duplex are set to auto-negotiate.

Configure Switch Ports and Power Over Ethernet

To configure switch ports and PoE, complete the following tasks.

Enable or Disable Switch Port Mode

You can set each interface independently to be either a firewall interface or a switch port. By default, Ethernet 1/1 is a firewall interface, and the remaining Ethernet interfaces are configured as switch ports.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Set the switch port mode by clicking the slider in the **SwitchPort** column so it shows as **Slider enabled** (🔵) or **Slider disabled** (⚫).

By default, switch ports are set to access mode in VLAN 1. You must manually add a logical VLAN 1 interface (or whichever VLAN you set for these switch ports) for traffic to be routed and to participate in the Firewall Threat Defense security policy (see Configure a VLAN Interface, on page 821). You cannot set the Management interface to switch port mode. When you change the switch port mode, all unsupported configuration is removed:

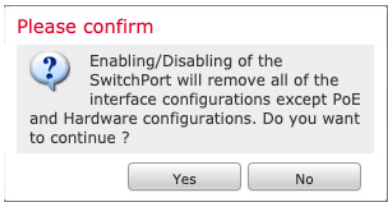

In addition, the interface is disabled if it was already enabled. Be sure to reenable the interface.

Configure a VLAN Interface

This section describes how to configure VLAN interfaces for use with associated switch ports. By default, switch ports are assigned to VLAN1; however, you must manually add the logical VLAN1 interface (or whichever VLAN you set for these switch ports) for traffic to be routed and to participate in the Firewall Threat Defense security policy.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Add Interfaces** > **VLAN Interface**.

Step 3 On **General**, set the following VLAN-specific parameters:

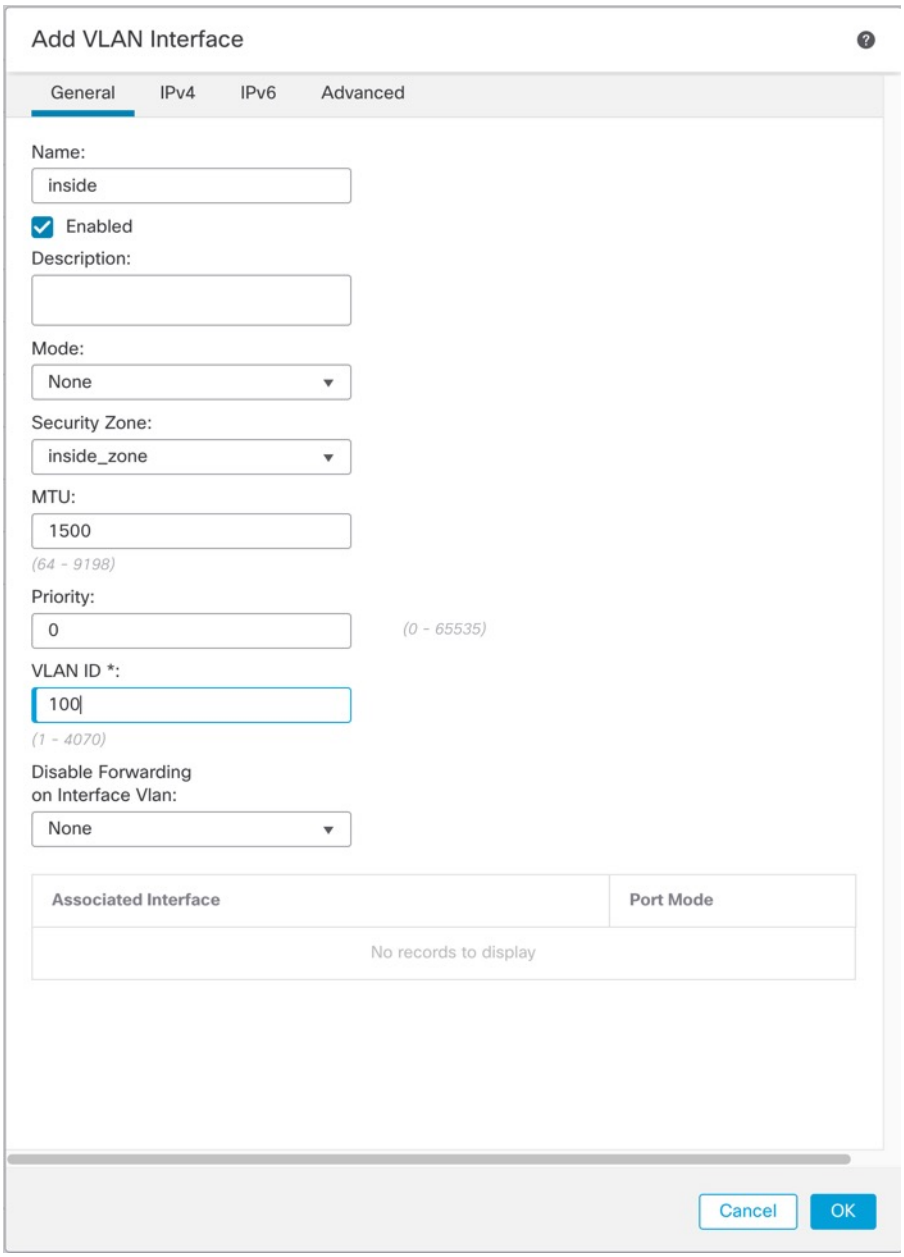

If you are editing an existing VLAN interface, the **Associated Interface** table shows switch ports on this VLAN.

a) Set the **VLAN ID**, between 1 and 4070, excluding IDs in the range 3968 to 4047, which are reserved for internal use.

   You cannot change the VLAN ID after you save the interface; the VLAN ID is both the VLAN tag used, and the interface ID in your configuration.

b) (Optional) Choose a VLAN ID for **Disable Forwarding on Interface VLAN** to disable forwarding to another VLAN.

For example, you have one VLAN assigned to the outside for internet access, one VLAN assigned to an inside business network, and a third VLAN assigned to your home network. The home network does not need to access the business network, so you can disable forwarding on the home VLAN; the business network can access the home network, but the home network cannot access the business network.

Step 4 To complete the interface configuration, see one of the following procedures:

- Configure Routed Mode Interfaces, on page 856
- Configure General Bridge Group Member Interface Parameters, on page 861

Step 5 Click **OK**.

Step 6 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Switch Ports as Access Ports

To assign a switch port to a single VLAN, configure it as an access port. Access ports accept only untagged traffic. The interfaces that are enabled with switch ports and assigned to VLAN1 are as follows:

Device Models	Switch Port Interfaces
Firepower 1010 and Secure Firewall 1210	Ethernet 1/2 through Ethernet 1/8
Secure Firewall 1220	Ethernet 1/2 through Ethernet 1/10

Note The device does not support Spanning Tree Protocol for loop detection in the network. Therefore you must ensure that any connection with the Firewall Threat Defense does not end up in a network loop.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Figure 359: Edit Physical Interface

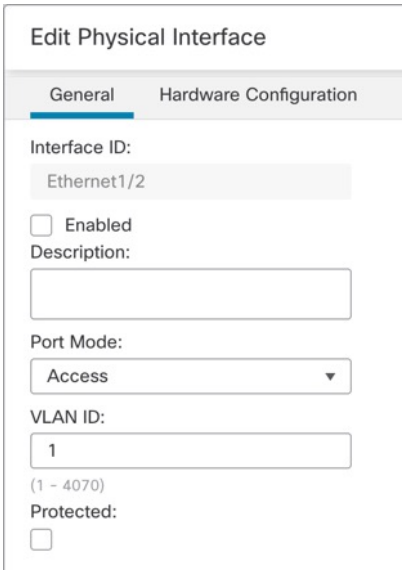

Step 3 Enable the interface by checking the **Enabled** check box.

Step 4 (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

Step 5 Set the **Port Mode** to **Access**.

Step 6 In the **VLAN ID** field, set the VLAN for this switch port, between 1 and 4070.

The default VLAN ID is 1.

Step 7 (Optional) Check the **Protected** check box to set this switch port as protected, so you can prevent the switch port from communicating with other protected switch ports on the same VLAN.

You might want to prevent switch ports from communicating with each other if: the devices on those switch ports are primarily accessed from other VLANs; you do not need to allow intra-VLAN access; and you want to isolate the devices from each other in case of infection or other security breach. For example, if you have a DMZ that hosts three web servers, you can isolate the web servers from each other if you enable **Protected** on each switch port. The inside and outside networks can both communicate with all three web servers, and vice versa, but the web servers cannot communicate with each other.

Step 8 (Optional) Set the duplex and speed by clicking **Hardware Configuration**.

Figure 360: Hardware Configuration

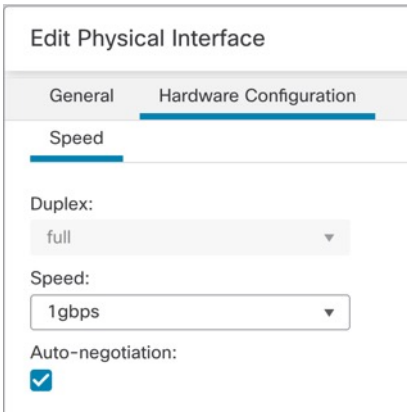

Check the **Auto-negotiation** check box (the default) to auto-detect the speed and duplex. If you uncheck it, you can set the speed and duplex manually:

- **Duplex**—Choose **Full** or **Half**.
- **Speed**—Choose **10mbps**, **100mbps**, or **1gbps**.

Step 9 Click **OK**.

Step 10 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Switch Ports as Trunk Ports

This procedure describes how to create a trunk port that can carry multiple VLANs using 802.1Q tagging. Trunk ports accept untagged and tagged traffic. Traffic on allowed VLANs pass through the trunk port unchanged.

When the trunk receives untagged traffic, it tags it to the native VLAN ID so that the Firewall Threat Defense device can forward the traffic to the correct switch ports, or can route it to another firewall interface. When the Firewall Threat Defense device sends native VLAN ID traffic out of the trunk port, it removes the VLAN tag. Be sure to set the same native VLAN on the trunk port on the other switch so that the untagged traffic will be tagged to the same VLAN.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Figure 361: Set Trunk Port Mode

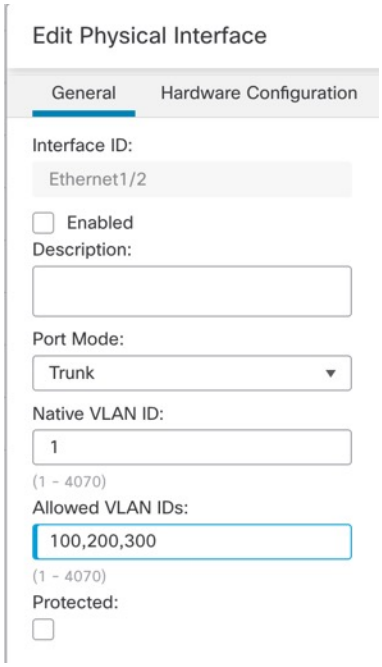

Step 3 Enable the interface by checking the **Enabled** check box.

Step 4 (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

Step 5 Set the **Port Mode** to **Trunk**.

Step 6 In the **Native VLAN ID** field, set the native VLAN for this switch port, between 1 and 4070.

The default native VLAN ID is 1.

Each port can only have one native VLAN, but every port can have either the same or a different native VLAN.

Step 7 In the **Allowed VLAN IDs** field, enter the VLANs for this trunk port between 1 and 4070.

You can identify up to 20 IDs in one of the following ways:

- A single number (n)
- A range (n-x)
- Numbers and ranges separated by commas, for example:

  5,7-10,13,45-100

  You can enter spaces instead of commas.

If you include the native VLAN in this field, it is ignored; the trunk port always removes the VLAN tagging when sending native VLAN traffic out of the port. Moreover, it will not receive traffic that still has native VLAN tagging.

Step 8 (Optional) Check the **Protected** check box to set this switch port as protected, so you can prevent the switch port from communicating with other protected switch ports on the same VLAN.

You might want to prevent switch ports from communicating with each other if: the devices on those switch ports are primarily accessed from other VLANs; you do not need to allow intra-VLAN access; and you want to isolate the devices from each other in case of infection or other security breach. For example, if you have a DMZ that hosts three web servers, you can isolate the web servers from each other if you enable **Protected** on each switch port. The inside and outside networks can both communicate with all three web servers, and vice versa, but the web servers cannot communicate with each other.

Step 9 (Optional) Set the duplex and speed by clicking **Hardware Configuration**.

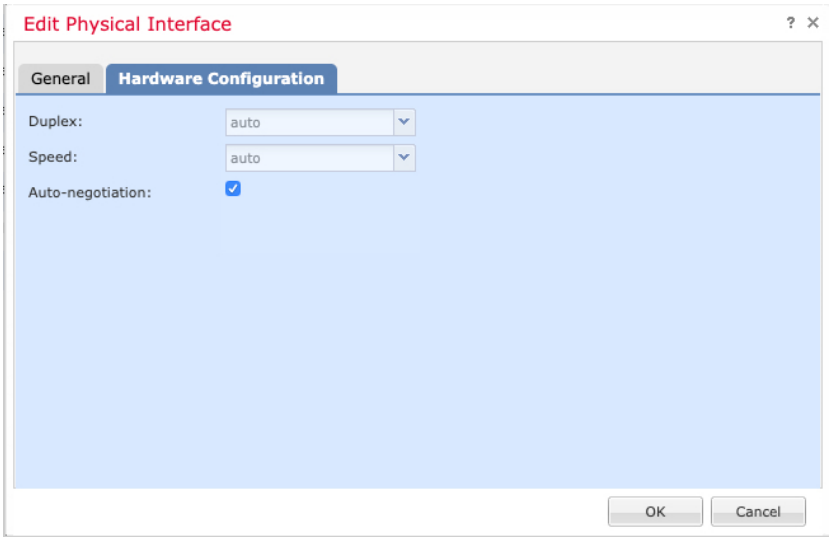

Check the **Auto-negotiation** check box (the default) to auto-detect the speed and duplex. If you uncheck it, you can set the speed and duplex manually:

- **Duplex**—Choose **Full** or **Half**.
- **Speed**—Choose **10mbps**, **100mbps**, or **1gbps**.

Step 10 Click **OK**.

Step 11 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Power Over Ethernet

Power over Ethernet (PoE) ports provide power for devices such as IP phones or wireless access points. PoE is enabled by default. This procedure describes how to disable and enable PoE and how to set optional parameters.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for Ethernet1/7 or 1/8 on Firepower 1010 or for any interface from Ethernet 1/5-1/8 on Secure Firewall 1210CP.

Step 3 Click **PoE**.

Figure 362: PoE

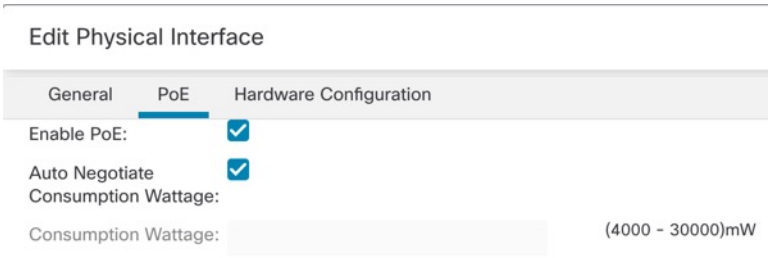

Step 4 Check the **Enable PoE** check box.

PoE is enabled by default.

Step 5 Choose autonegotiation or manual power.

- **Auto Negotiate Consumption Wattage**—Delivers power automatically to the powered device using a wattage appropriate to the class of the powered device. The firewall uses LLDP to further negotiate the correct wattage. When you connect a device of a certain class, it will be provisioned up to the maximum for that class in case it ever needs to use more power. For example, if you add a class 4 device that requests 12.95W, it will be allocated 30W even if it doesn't currently use all that power. Some devices can renegotiate power needs. If you know your device needs less power than is allocated, you can instead set the **Consumption Wattage** manually to free up the power for other devices.
- **Consumption Wattage**—Uncheck the **Auto Negotiate Consumption Wattage** check box to set the manual wattage to manually specify the wattage in milliwatts, from 4000 to 30000 (1010) or 90000 (1210CP). Use this option if you want to set the watts manually and disable LLDP negotiation. For manual allocation, the class will show in **show power inline** output as **n/a** because the class is not used to decide the power consumption.

  View the current PoE status using the **show power inline** command.

Step 6 Click **OK**.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Loopback Interfaces

This section tells how to configure loopback interfaces.

About Loopback Interfaces

A loopback interface is a software-only interface that emulates a physical interface. This interface is reachable on IPv4 and IPv6 through multiple physical interfaces. The loopback interface helps to overcome path failures; it is accessible from any physical interface, so if one goes down, you can access the loopback interface from another.

Loopback interfaces can be used for:

- AAA
- BGP
- DNS
- HTTP
- ICMP
- IPsec flow offload—Secure Firewall 1200, 3100 and 4200 only
- NetFlow
- SNMP
- SSH
- Static and dynamic VTI tunnels
- Syslog

The Firewall Threat Defense can distribute the loopback address using dynamic routing protocols, or you can configure a static route on the peer device to reach the loopback IP address through one of the Firewall Threat Defense's physical interfaces. You cannot configure a static route on the Firewall Threat Defense that specifies the loopback interface.

Related Topics

Guidelines and Limitations for Loopback Interfaces, on page 829
Configure a Loopback Interface, on page 830

Guidelines and Limitations for Loopback Interfaces

Firewall Mode

- Supported in routed mode only.

High Availability and Clustering

- No clustering support.

Additional Guidelines and Limitations

- TCP sequence randomization is always disabled for traffic from the physical interface to the loopback interface.

Configure a Loopback Interface

To add a loopback interface for a device:

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 From the **Add Interfaces** drop-down list, choose **Loopback Interface**.

Step 3 In the **General** tab, configure the following parameters:

a) **Name**—Enter a name for the loopback interface.
b) **Enabled**—Check the check box to enable the loopback interface.
c) **Loopback ID**—Enter the loopback ID between 1 to 1024.
d) **Description**—Enter a description for the loopback interface.

Step 4 Configure the routed mode interface parameters. See Configure Routed Mode Interfaces, on page 856.

Rate-Limit Traffic to the Loopback Interface

Before you begin

You should rate-limit traffic going to the loopback interface IP address to prevent excessive load on the system. You can add a connection limit rule to the global service policy.

Procedure

Step 1 Create an extended access list identifying traffic to the loopback interface IP address(es).

a) Choose **Objects** > **Object Management** > **Access List** > **Extended**.
b) Click **Add Extended Access List** to create a new ACL.
c) In the **New Extended Access List Object** dialog box, enter a name for the ACL (no spaces allowed), and click **Add** to create a new entry.

Figure 363: Name ACL and Add Entry

d) Configure the source (any) and destination addresses (loopback IP addresses) on the **Network** tab.

Figure 364: Source and Destination Networks

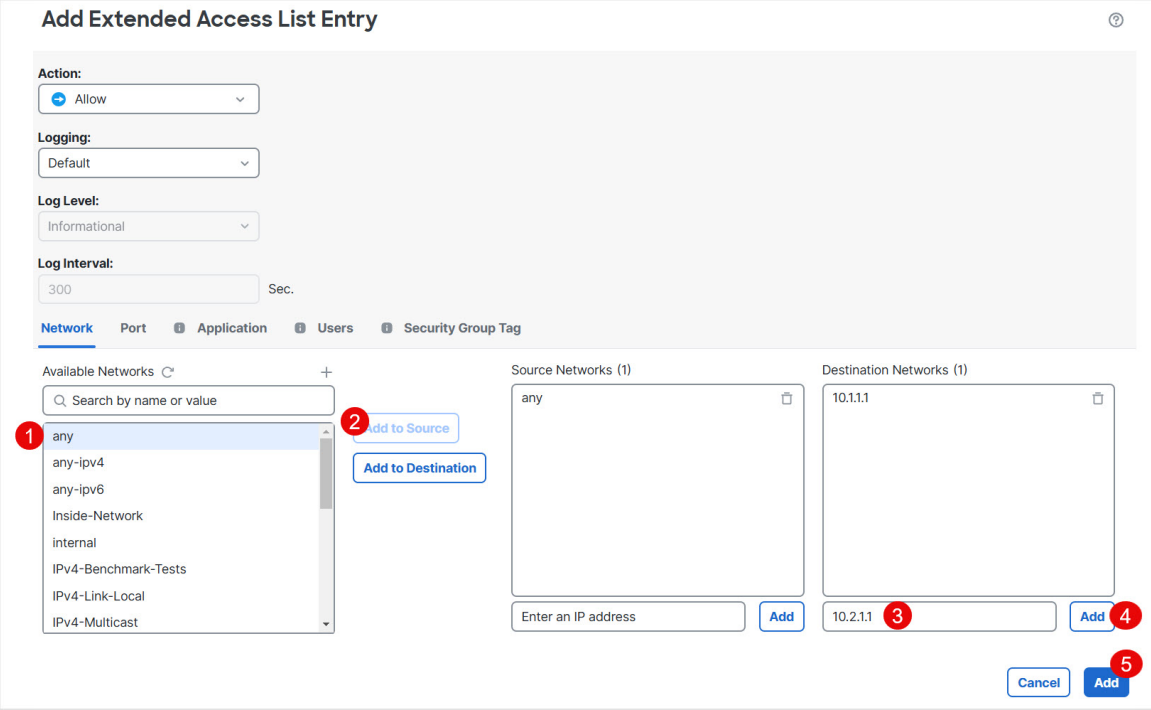

Note
Keep the default **Action** as **Allow** (match) and other settings as-is.

- Source—Select **any** from the **Available Networks** list, and click **Add to Source**. You can also narrow this access list by specifying the source IP addresses instead of **any**.
- Destination—Type an address in the edit box below the **Destination Networks** list and click **Add**. Repeat for each loopback interface.

e) Click **Add** to add the entry to the ACL.
f) Click **Save** to save the ACL.

Figure 365: Save ACL

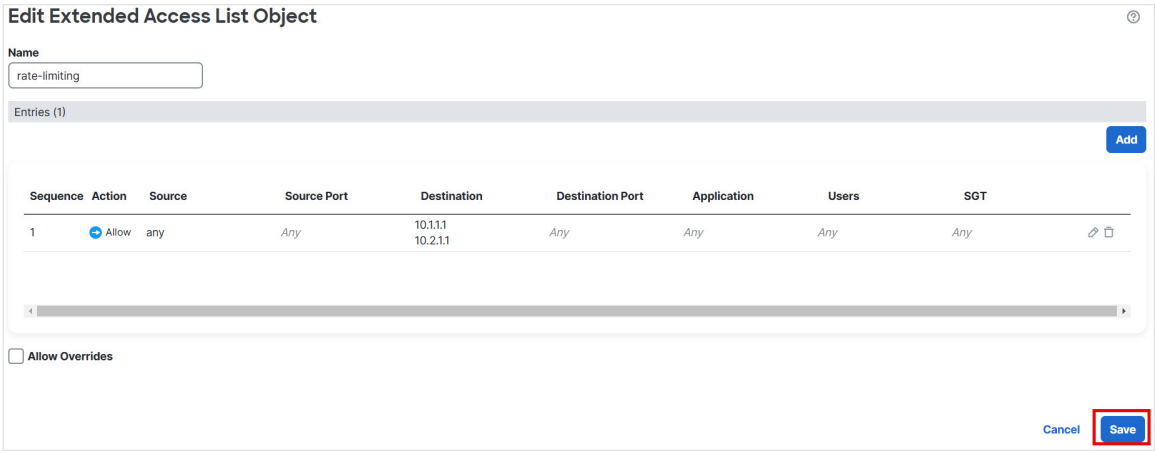

Step 2 Choose **Policies** > **Access Control heading** > **Access Control**, and click **Edit** (✎) for the access control policy assigned to your device.

Step 3 Click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

Figure 366: Advanced Settings

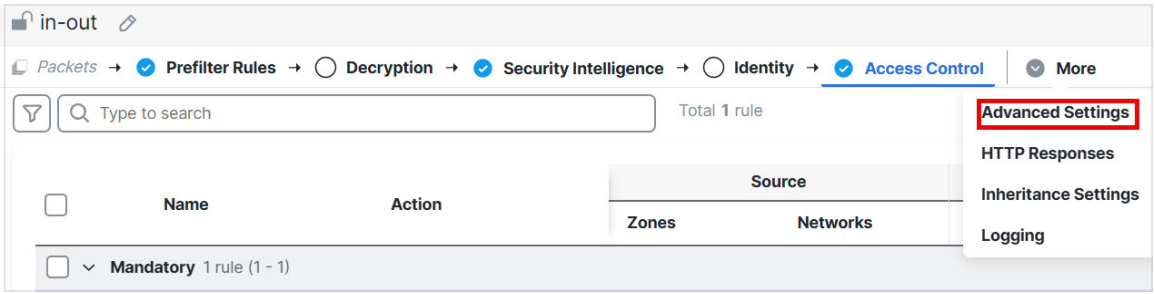

Step 4 Click **Edit** (✎) in the **Threat Defense Service Policy** group.

Figure 367: Threat Defense Service Policy

Step 5 Click **Add Rule** to create a new rule.

Figure 368: Add Rule

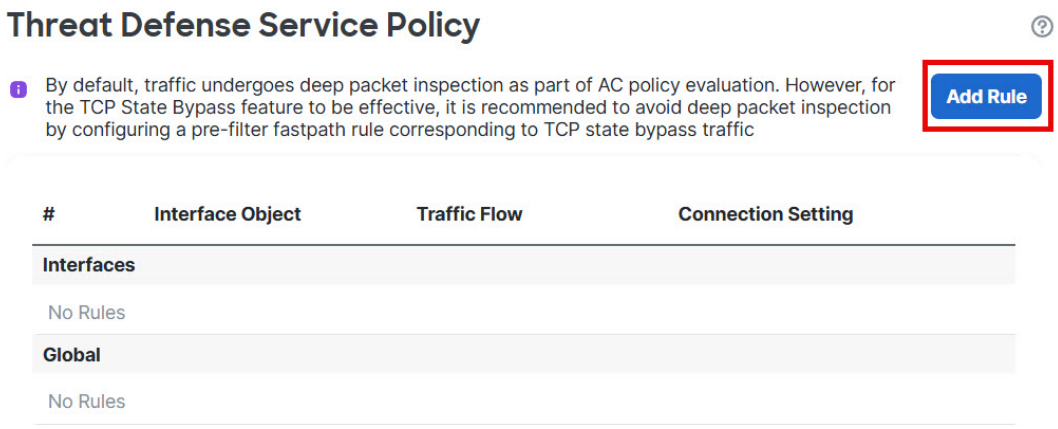

The service policy rule wizard opens to step you through the process of configuring the rule.

Step 6 In the **Interface Object** step, click **Global** to create a global rule, which applies to all interfaces, then click **Next**.

Figure 369: Global Policy

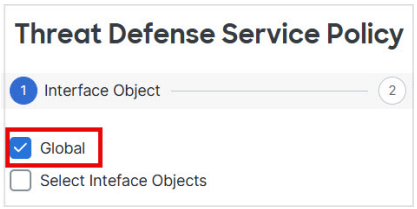

Step 7 In the **Traffic Flow** step, select the extended access list object you created in Step 1, on page 830, and then click **Next**.

Figure 370: Choose Extended Access List

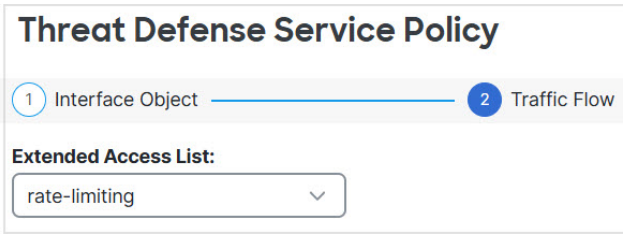

Step 8 In the **Connection Setting** step, set the **Connections** limits.

Figure 371: Set Connection Limits

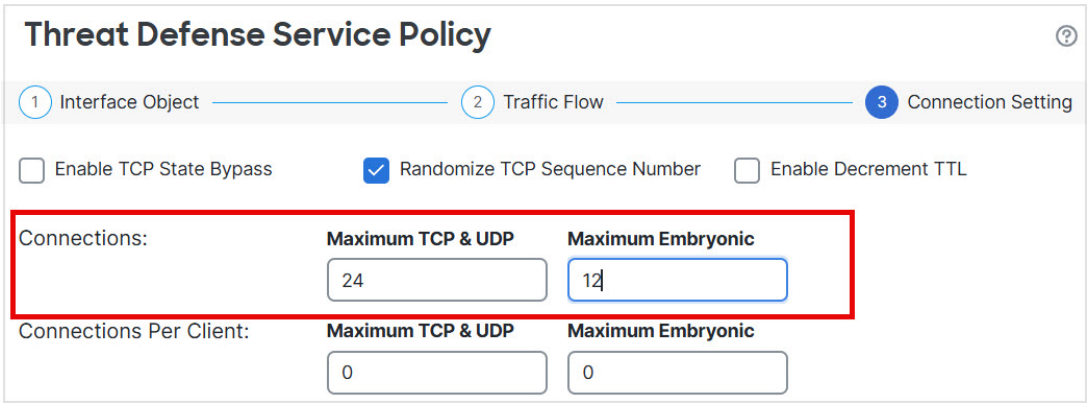

Set the **Maximum TCP & UDP** connections to the expected number of connections for the loopback interface, and the **Maximum Embryonic** connections to a lower number. For example, you can set it to 5/2, or 10/5, or 1024/512, depending on the expected loopback interface sessions you need.

Setting the embryonic connection limit enables TCP Intercept, which protects the system from a DoS attack perpetrated by flooding an interface with TCP SYN packets.

Step 9 Click **Finish** to save your changes.

Step 10 Click **OK**.

Step 11 Click **Save** on the **Advanced Settings** window.

Step 12 You can now deploy the changes to the affected devices.

Configure VLAN Subinterfaces and 802.1Q Trunking

VLAN subinterfaces let you divide a physical, redundant, or EtherChannel interface into multiple logical interfaces that are tagged with different VLAN IDs. An interface with one or more VLAN subinterfaces is automatically configured as an 802.1Q trunk. Because VLANs let you keep traffic separate on a given physical interface, you can increase the number of interfaces available to your network without adding additional physical interfaces or devices.

Guidelines and Limitations for VLAN Subinterfaces

Model Support

- Firepower 1010 and Secure Firewall 1210/1220—VLAN subinterfaces are not supported on switch ports or VLAN interfaces.
- Firepower 1010 and Secure Firewall 1210/1220—You cannot create a subinterface using interface ID 1 or VLAN ID 1. VLAN 1 is reserved for the logical VLAN interface for switch ports.

High Availability and Clustering

You cannot use a subinterface for the failover or state link or for the cluster control link. The exception is for multi-instance mode: you can use a *chassis*-defined subinterface for these links.

Additional Guidelines

- Preventing untagged packets on the physical interface—If you use subinterfaces, you typically do not also want the physical interface to pass traffic, because the physical interface passes untagged packets. This property is also true for the active physical interface in a redundant interface pair and for EtherChannel links. Because the physical, redundant, or EtherChannel interface must be enabled for the subinterface to pass traffic, ensure that the physical, redundant, or EtherChannel interface does not pass traffic by not configuring a name for the interface. If you want to let the physical, redundant, or EtherChannel interface pass untagged packets, you can configure the name as usual.
- You cannot configure subinterfaces on the Management interface, either the dedicated Management interface configured at the CLI nor a data interface used for manager access.
- All subinterfaces on the same parent interface must be either bridge group members or routed interfaces; you cannot mix and match.
- The Firewall Threat Defense does not support the Dynamic Trunking Protocol (DTP), so you must configure the connected switch port to trunk unconditionally.
- You might want to assign unique MAC addresses to subinterfaces defined on the Firewall Threat Defense, because they use the same burned-in MAC address of the parent interface. For example, your service provider might perform access control based on the MAC address. Also, because IPv6 link-local addresses are generated based on the MAC address, assigning unique MAC addresses to subinterfaces allows for unique IPv6 link-local addresses, which can avoid traffic disruption in certain instances on the Firewall Threat Defense.

Note If you manually assign a MAC address, be sure to assign MAC addresses to all subinterfaces on the same physical interface to avoid unexpected behavior and outages.

Maximum Number of VLAN Subinterfaces by Device Model

The device model limits the maximum number of VLAN subinterfaces that you can configure. Note that you can configure subinterfaces on data interfaces only, you cannot configure them on the management interface.

The following table explains the limits for each device model.

Model	Maximum VLAN Subinterfaces
Firepower 1010	60
Firepower 1120	512
Firepower 1140, 1150	1024
Secure Firewall 1200	1024

Model	Maximum VLAN Subinterfaces
Secure Firewall 3100, 4200	1024
Firepower 4100	1024
Firepower 9300	1024
Firewall Threat Defense Virtual	50
ISA 3000	100

Add a Subinterface

Add one or more subinterfaces to a physical, redundant, or port-channel interface.

For the Firepower 4100/9300, you can configure subinterfaces in FXOS for use with container instances; see Add a VLAN Subinterface for Container Instances, on page 335. These subinterfaces appear in the Firewall Management Center interface list. You can also add subinterfaces in Firewall Management Center, but only on parent interfaces that do not already have subinterfaces defined in FXOS.

Note The parent physical interface passes untagged packets. You may not want to pass untagged packets, so be sure not to include the parent interface in your security policy.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Enable the parent interface according to Enable the Physical Interface and Configure Ethernet Settings, on page 778.

Step 3 Click **Add Interfaces** > **Sub Interface**.

Step 4 On **General**, set the following parameters:

Figure 372: Add Subinterface

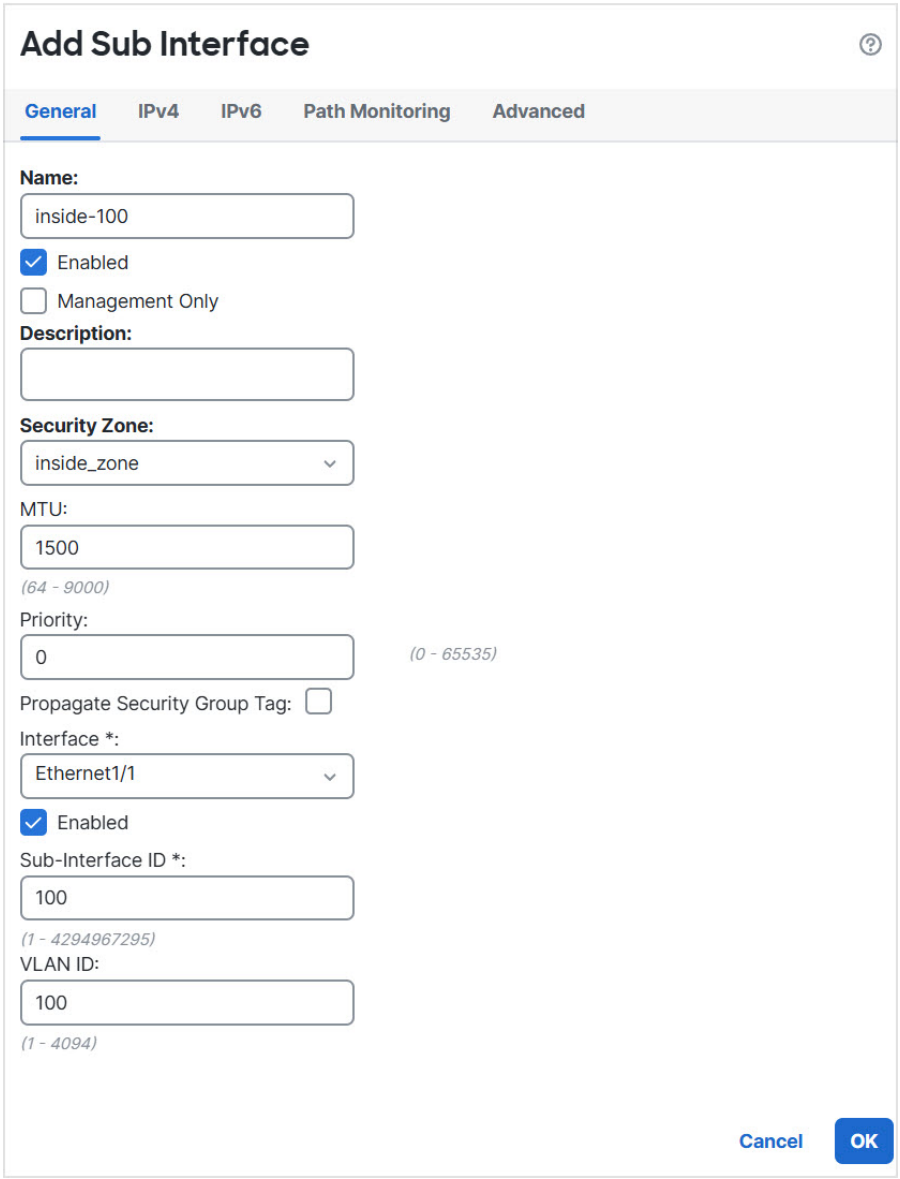

a) **Interface**—Choose the physical, redundant, or port-channel interface to which you want to add the subinterface.

b) **Sub-Interface ID**—Enter the subinterface ID as an integer between 1 and 4294967295. The number of subinterfaces allowed depends on your platform. You cannot change the ID after you set it.

c) **VLAN ID**—Enter the VLAN ID between 1 and 4094 that will be used to tag the packets on this subinterface.

   This VLAN ID must be unique.

Step 5 Click **OK**.

Step 6 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Step 7 Configure the routed or transparent mode interface parameters. See Configure Routed Mode Interfaces, on page 856 or Configure Bridge Group Interfaces, on page 861.

Configure VXLAN Interfaces

This chapter tells how to configure Virtual eXtensible LAN (VXLAN) interfaces. VXLAN interfaces act as Layer 2 virtual networks over Layer 3 physical networks to stretch Layer 2 networks.

About VXLAN Interfaces

VXLAN provides the same Ethernet Layer 2 network services as VLAN does, but with greater extensibility and flexibility. Compared to VLAN, VXLAN offers the following benefits:

- Flexible placement of multitenant segments throughout the data center.
- Higher scalability to address more Layer 2 segments: up to 16 million VXLAN segments.

This section describes how VXLAN works. For detailed information about VXLAN, see RFC 7348. For detailed information about Geneve, see RFC 8926.

Encapsulation

The Firewall Threat Defense supports two types of VXLAN encapsulation:

- VXLAN (all models)—VXLAN uses MAC Address-in-User Datagram Protocol (MAC-in-UDP) encapsulation. The original Layer 2 frame has a VXLAN header added and is then placed in a UDP-IP packet.
- Geneve (Firewall Threat Defense Virtual only)—Geneve has a flexible inner header that is not limited to the MAC address. Geneve encapsulation is required for transparent routing of packets between an Amazon Web Services (AWS) Gateway Load Balancer and appliances, and for sending extra information.

VXLAN Tunnel Endpoint

VXLAN tunnel endpoint (VTEP) devices perform VXLAN encapsulation and decapsulation. Each VTEP has two interface types: one or more virtual interfaces called VXLAN Network Identifier (VNI) interfaces to which you apply your security policy, and a regular interface called the VTEP source interface that tunnels the VNI interfaces between VTEPs. The VTEP source interface is attached to the transport IP network for VTEP-to-VTEP communication.

The following figure shows two Firewall Threat Defenses and Virtual Server 2 acting as VTEPs across a Layer 3 network, extending the VNI 1, 2, and 3 networks between sites. The Firewall Threat Defenses act as bridges or gateways between VXLAN and non-VXLAN networks.

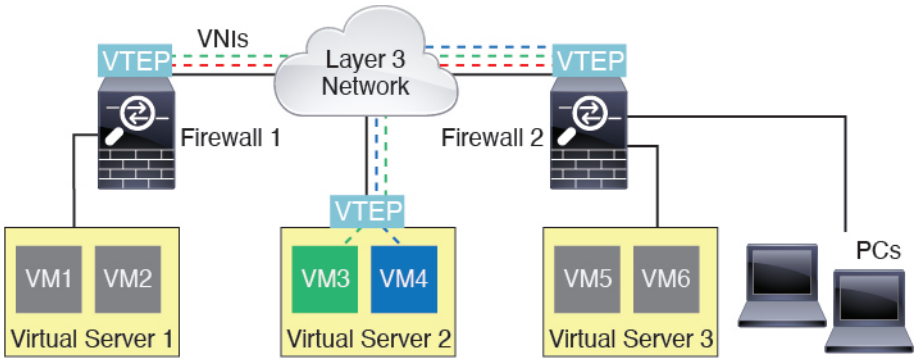

The underlying IP network between VTEPs is independent of the VXLAN overlay. Encapsulated packets are routed based on the outer IP address header, which has the initiating VTEP as the source IP address and the terminating VTEP as the destination IP address. For VXLAN encapsulation: The destination IP address can be a multicast group when the remote VTEP is not known. With Geneve, the Firewall Threat Defense only supports static peers. The destination port for VXLAN is UDP port 4789 by default (user configurable). The destination port for Geneve is 6081.

VTEP Source Interface

The VTEP source interface is a regular interface (physical, EtherChannel, or even VLAN) with which you plan to associate all VNI interfaces. You can configure one VTEP source interface per Firewall Threat Defense Virtual. Because you can only configure one VTEP source interface, you cannot configure both VXLAN and Geneve interfaces on the same device. There is an exception for Firewall Threat Defense Virtual clustering on AWS or Azure, where you can have two VTEP source interfaces: a VXLAN interface is used for the cluster control link, and a Geneve (AWS) or VXLAN (Azure) interface can be used for the Gateway Load Balancer.

The VTEP source interface can be devoted wholly to VXLAN traffic, although it is not restricted to that use. If desired, you can use the interface for regular traffic and apply a security policy to the interface for that traffic. For VXLAN traffic, however, all security policy must be applied to the VNI interfaces. The VTEP interface serves as a physical port only.

In transparent firewall mode, the VTEP source interface is not part of a BVI, and you do configure an IP address for it, similar to the way the management interface is treated.

VNI Interfaces

VNI interfaces are similar to VLAN interfaces: they are virtual interfaces that keep network traffic separated on a given physical interface by using tagging. You apply your security policy directly to each VNI interface.

You can only add one VTEP interface, and all VNI interfaces are associated with the same VTEP interface. There is an exception for Firewall Threat Defense Virtual clustering on AWS or Azure. For AWS clustering, you can have two VTEP source interfaces: a VXLAN interface is used for the cluster control link, and a Geneve interface can be used for the AWS Gateway Load Balancer. For Azure clustering, you can have two VTEP source interfaces: a VXLAN interface is used for the cluster control link, and a second VXLAN interface can be used for the Azure Gateway Load Balancer.

VXLAN Packet Processing

VXLAN

Traffic entering and exiting the VTEP source interface is subject to VXLAN processing, specifically encapsulation or decapsulation.

Encapsulation processing includes the following tasks:

- The VTEP source interface encapsulates the inner MAC frame with the VXLAN header.
- The UDP checksum field is set to zero.
- The Outer frame source IP is set to the VTEP interface IP.
- The Outer frame destination IP is decided by a remote VTEP IP lookup.

Decapsulation; the Firewall Threat Defense only decapsulates a VXLAN packet if:

- It is a UDP packet with the destination port set to 4789 (this value is user configurable).
- The ingress interface is the VTEP source interface.
- The ingress interface IP address is the same as the destination IP address.
- The VXLAN packet format is compliant with the standard.

Geneve

Traffic entering and exiting the VTEP source interface is subject to Geneve processing, specifically encapsulation or decapsulation.

Encapsulation processing includes the following tasks:

- The VTEP source interface encapsulates the inner MAC frame with the Geneve header.
- The UDP checksum field is set to zero.
- The Outer frame source IP is set to the VTEP interface IP.
- The Outer frame destination IP is set the peer IP address that you configured.

Decapsulation; the Firewall Threat Defense only decapsulates a Geneve packet if:

- It is a UDP packet with the destination port set to 6081 (this value is user configurable).
- The ingress interface is the VTEP source interface.
- The ingress interface IP address is the same as the destination IP address.
- The Geneve packet format is compliant with the standard.

Peer VTEPs

When the Firewall Threat Defense sends a packet to a device behind a peer VTEP, the Firewall Threat Defense needs two important pieces of information:

- The destination MAC address of the remote device

- The destination IP address of the peer VTEP

The Firewall Threat Defense maintains a mapping of destination MAC addresses to remote VTEP IP addresses for the VNI interfaces.

VXLAN Peer

There are two ways in which the Firewall Threat Defense can find this information:

- A single peer VTEP IP address can be statically configured on the Firewall Threat Defense.

  For IPv4: The Firewall Threat Defense then sends a VXLAN-encapsulated ARP broadcast to the VTEP to learn the end node MAC address.

  For IPv6: The Firewall Threat Defense then sends an IPv6 Neighbor Solicitation message to the IPv6 solicited-node multicast address. The peer VTEP responds with an IPv6 Neighbor Advertisement message with its link-local address.

- A group of peer VTEP IP addresses can be statically configured on the Firewall Threat Defense.

  For IPv4: The Firewall Threat Defense then sends a VXLAN-encapsulated ARP broadcast to the VTEP to learn the end node MAC addresses.

  For IPv6: The Firewall Threat Defense then sends an IPv6 Neighbor Solicitation message to the IPv6 solicited-node multicast address. The peer VTEP responds with an IPv6 Neighbor Advertisement message with its link-local address.

- A multicast group can be configured on each VNI interface (or on the VTEP as a whole).

  For IPv4: The Firewall Threat Defense sends a VXLAN-encapsulated ARP broadcast packet within an IP multicast packet through the VTEP source interface. The response to this ARP request enables the Firewall Threat Defense to learn both the remote VTEP IP address along with the destination MAC address of the remote end node.

  For IPv6: The Firewall Threat Defense sends a Multicast Listener Discovery (MLD) Report message through the VTEP source interface to indicate that the Firewall Threat Defense is listening on the VTEP interface for the multicast address traffic.

  This option is not supported with Geneve.

Geneve Peer

The Firewall Threat Defense Virtual only supports statically defined peers. You can define the Firewall Threat Defense Virtual peer IP address on the AWS Gateway Load Balancer. Because the Firewall Threat Defense Virtual never initiates traffic to the Gateway Load Balancer, you do not also have to specify the Gateway Load Balancer IP address on the Firewall Threat Defense Virtual; it learns the peer IP address when it receives Geneve traffic. Multicast groups are not supported with Geneve.

VXLAN Use Cases

This section describes the use cases for implementing VXLAN on the Firewall Threat Defense.

VXLAN Bridge or Gateway Overview

Each Firewall Threat Defense VTEP acts as a bridge or gateway between end nodes such as VMs, servers, and PCs and the VXLAN overlay network. For incoming frames received with VXLAN encapsulation over the VTEP source interface, the Firewall Threat Defense strips out the VXLAN header and forwards it to a

physical interface connected to a non-VXLAN network based on the destination MAC address of the inner Ethernet frame.

The Firewall Threat Defense always processes VXLAN packets; it does not just forward VXLAN packets untouched between two other VTEPs.

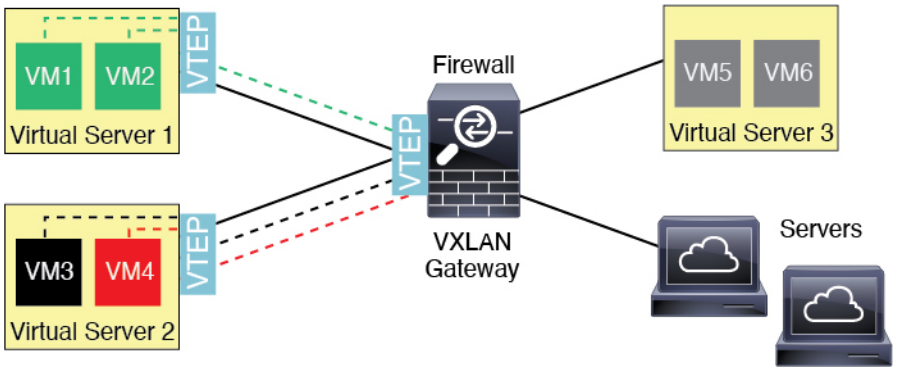

VXLAN Bridge

When you use a bridge group (transparent firewall mode, or optionally routed mode), the Firewall Threat Defense can serve as a VXLAN bridge between a (remote) VXLAN segment and a local segment where both are in the same network. In this case, one member of the bridge group is a regular interface while the other member is a VNI interface.

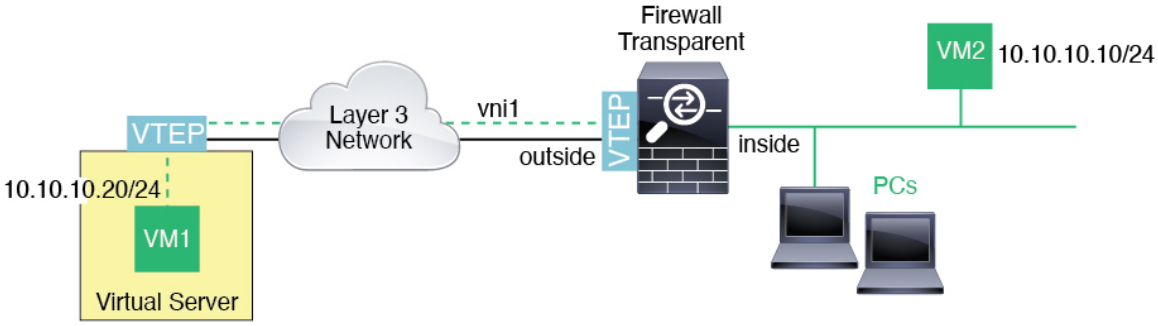

VXLAN Gateway (Routed Mode)

The Firewall Threat Defense can serve as a router between VXLAN and non-VXLAN domains, connecting devices on different networks.

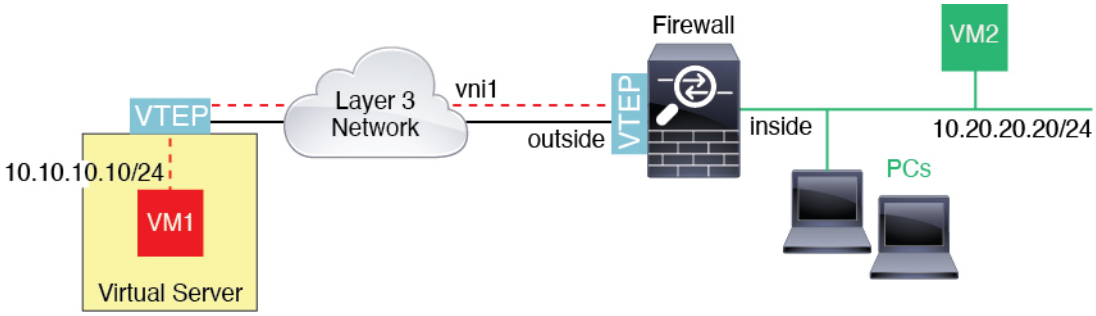

Router Between VXLAN Domains

With a VXLAN-stretched Layer 2 domain, a VM can point to an Firewall Threat Defense as its gateway while the Firewall Threat Defense is not on the same rack, or even when the Firewall Threat Defense is far away over the Layer 3 network.

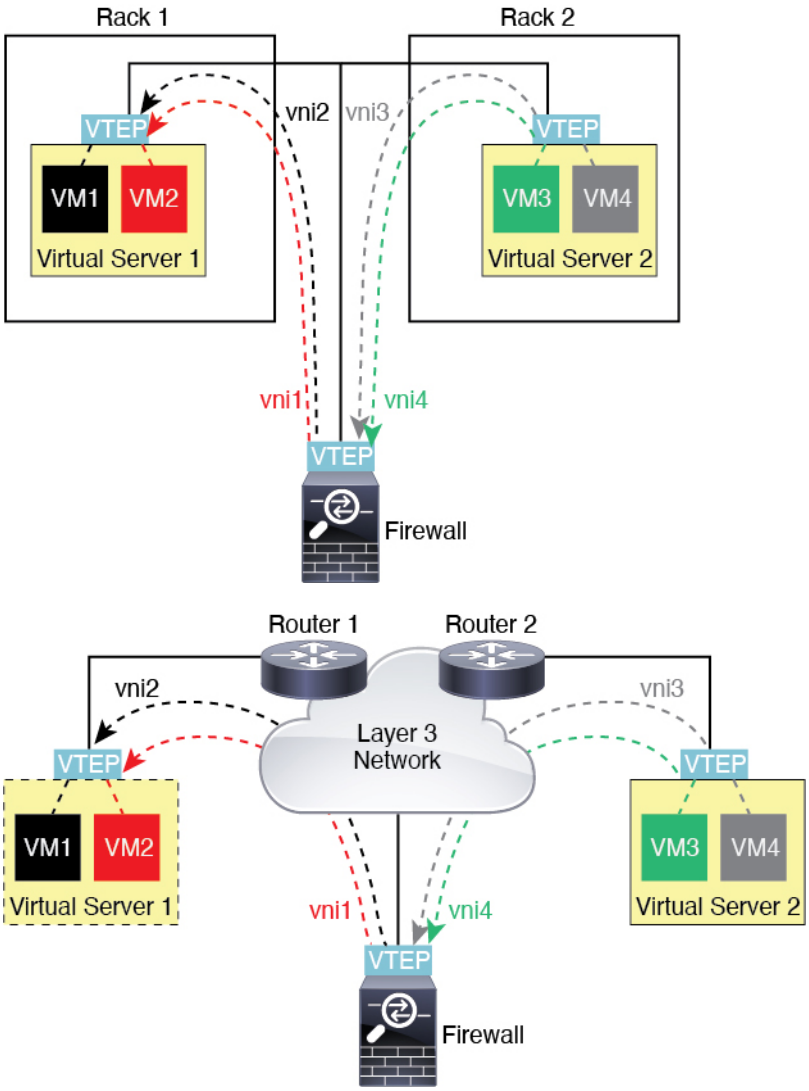

See the following notes about this scenario:

1. For packets from VM3 to VM1, the destination MAC address is the Firewall Threat Defense MAC address, because the Firewall Threat Defense is the default gateway.
2. The VTEP source interface on Virtual Server 2 receives packets from VM3, then encapsulates the packets with VNI 3's VXLAN tag and sends them to the Firewall Threat Defense.
3. When the Firewall Threat Defense receives the packets, it decapsulates the packets to get the inner frames.
4. The Firewall Threat Defense uses the inner frames for route lookup, then finds that the destination is on VNI 2. If it does not already have a mapping for VM1, the Firewall Threat Defense sends an encapsulated ARP broadcast on the multicast group IP on VNI 2.

Note The Firewall Threat Defense must use dynamic VTEP peer discovery because it has multiple VTEP peers in this scenario.

5. The Firewall Threat Defense encapsulates the packets again with the VXLAN tag for VNI 2 and sends the packets to Virtual Server 1. Before encapsulation, the Firewall Threat Defense changes the inner frame destination MAC address to be the MAC of VM1 (multicast-encapsulated ARP might be needed for the Firewall Threat Defense to learn the VM1 MAC address).
6. When Virtual Server 1 receives the VXLAN packets, it decapsulates the packets and delivers the inner frames to VM1.

Geneve Single-Arm Proxy

Note This use case is the only currently supported use case for Geneve interfaces.

The AWS Gateway Load Balancer combines a transparent network gateway and a load balancer that distributes traffic and scales virtual appliances on demand. The Firewall Threat Defense Virtual supports the Gateway Load Balancer centralized control plane with a distributed data plane (Gateway Load Balancer endpoint). The following figure shows traffic forwarded to the Gateway Load Balancer from the Gateway Load Balancer endpoint. The Gateway Load Balancer balances traffic among multiple Firewall Threat Defense Virtuals, which inspect the traffic before either dropping it or sending it back to the Gateway Load Balancer (U-turn traffic). The Gateway Load Balancer then sends the traffic back to the Gateway Load Balancer endpoint and to the destination.

Figure 373: Geneve Single-Arm Proxy

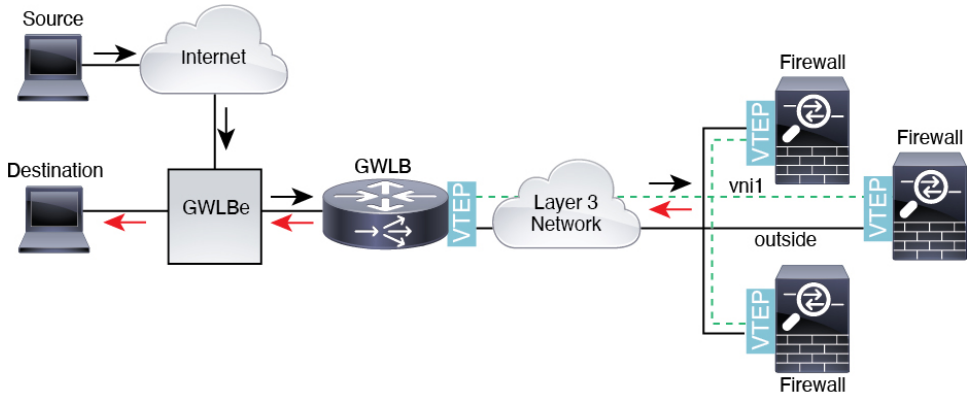

AWS Gateway Load Balancer and Geneve Dual-Arm Proxy

Note This use case is the only supported use case for Geneve interfaces.

The Firewall Threat Defense Virtual supports the Gateway Load Balancer centralized control plane with a distributed data plane (Gateway Load Balancer endpoint) in single-arm and dual-arm mode. The following figure shows outbound traffic (traffic inspected by Firewall Threat Defense Virtual) directly forwarded to the

destination (internet) without the need for traffic hop to the GWLB and GWLB endpoint. The Firewall Threat Defense Virtual inspect the outbound traffic and perform NAT before either dropping it or sending it back to the internet through NAT gateway. Dual-arm proxy provides a common egress path for multi-VPC deployment. The firewall inspects the outbound traffic from multiple VPCs, and it exits from a single point to the internet, making it a cost-effective infrastructure solution.

Figure 374: Geneve Dual-Arm Proxy: Egress Traffic from Single VPC

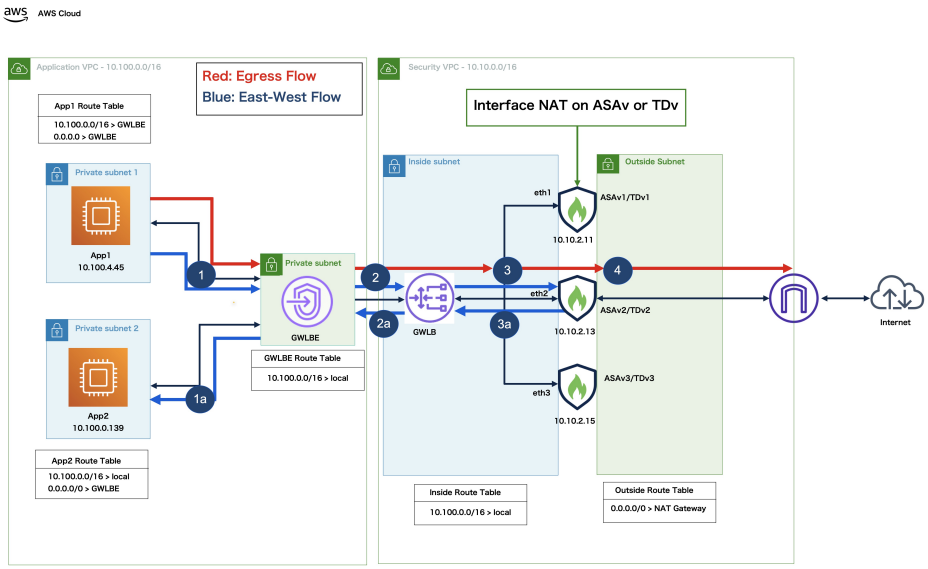

Figure 375: Geneve Dual-Arm Proxy: Egress Traffic from Multiple VPCs

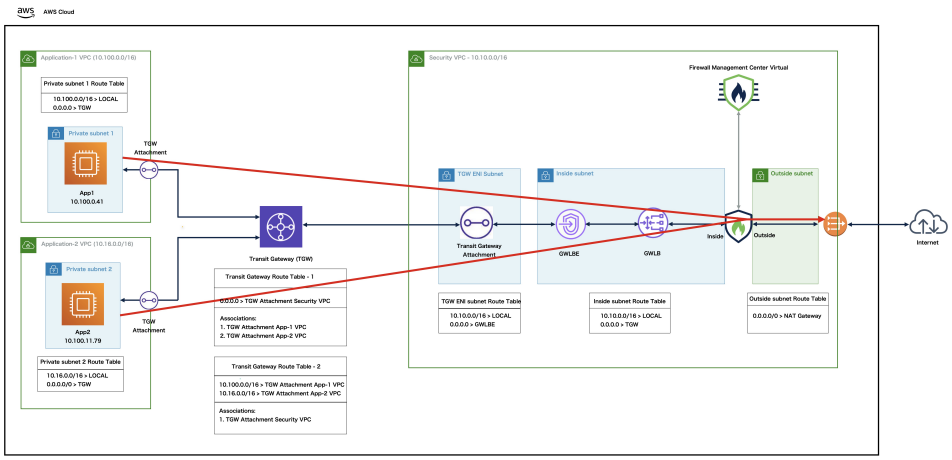

Azure Gateway Load Balancer and Paired Proxy

In an Azure service chain, Firewall Threat Defense Virtuals act as a transparent gateway that can intercept packets between the internet and the customer service. The Firewall Threat Defense Virtual defines an external interface and an internal interface on a single NIC by utilizing VXLAN segments in a paired proxy.

The following figure shows traffic forwarded to the Azure Gateway Load Balancer from the Public Gateway Load Balancer on the external VXLAN segment. The Gateway Load Balancer balances traffic among multiple

Firewall Threat Defense Virtuals, which inspect the traffic before either dropping it or sending it back to the Gateway Load Balancer on the internal VXLAN segment. The Azure Gateway Load Balancer then sends the traffic back to the Public Gateway Load Balancer and to the destination.

Figure 376: Azure Gateway Load Balancer with Paired Proxy

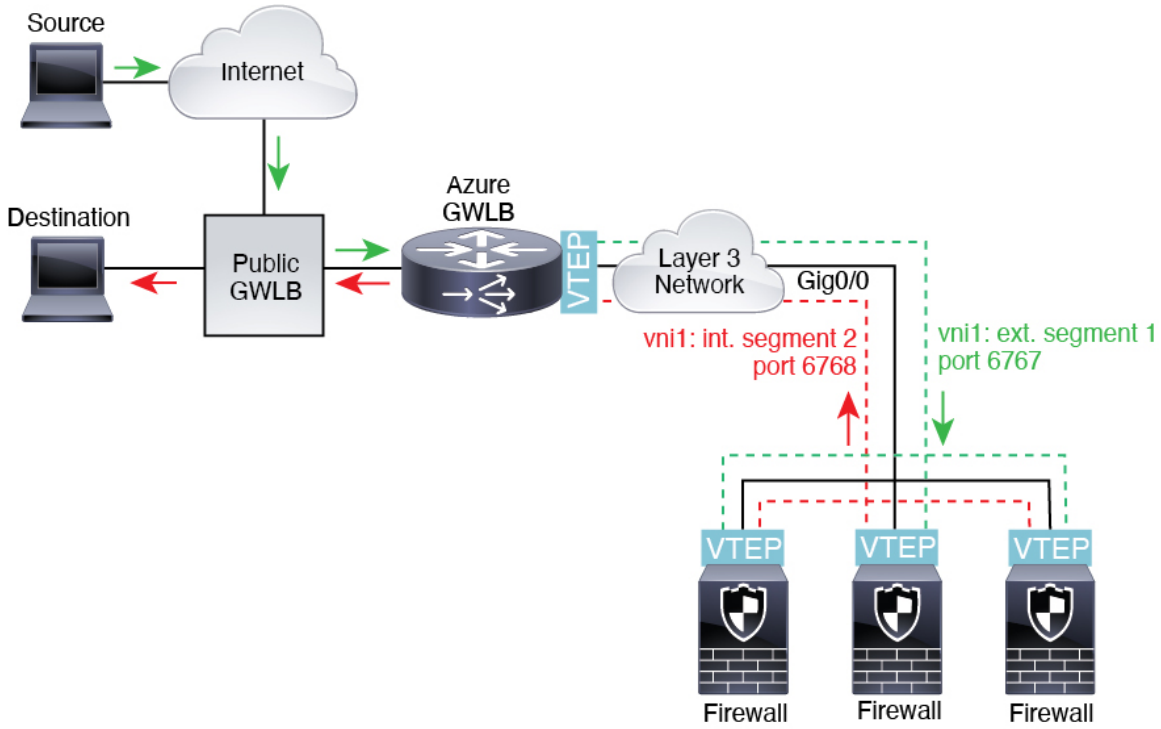

Requirements and Prerequisites for VXLAN Interfaces

Model Requirements

- VXLAN encapsulation is supported on all models.
- Geneve encapsulation is supported for the following models:
  - Firewall Threat Defense Virtual in Amazon Web Services (AWS)
- VXLAN in *paired proxy mode* is supported for the following models:
  - Firewall Threat Defense Virtual in Azure
- Firepower 1010 and Secure Firewall 1210/1220 switch ports and VLAN interfaces are not supported as VTEP interfaces.

Guidelines for VXLAN Interfaces

Firewall Mode

- Geneve interfaces are only supported in routed firewall mode.
- Paired proxy VXLAN interfaces are only supported in routed firewall mode.

IPv6

- The VNI interface supports both IPv4 and IPv6 traffic.
- For VXLAN encapsulation, the VTEP source interface supports both IPv4 and IPv6. The Firewall Threat Defense Virtual cluster control link VTEP source interface only supports IPv4.

  For Geneve, the VTEP source interface only supports IPv4.

Clustering

- Clustering does not support VXLAN in Individual Interface mode except for the cluster control link (Firewall Threat Defense Virtual only) Only Spanned EtherChannel mode supports VXLAN.

  An exception is made for AWS, which can use an additional Geneve interface for use with the GWLB and for Azure, which can use an additional paired proxy VXLAN interface for use with the GWLB.

Routing

- Only static routing or Policy Based Routing is supported on the VNI interface; dynamic routing protocols are not supported.

VPN

You cannot configure the VTEP source interface for VPN or use it as a VTI.

MTU

- VXLAN encapsulation—If the source interface MTU is less than 1554 bytes for IPv4 or 1574 bytes for IPv6, then the Firewall Threat Defense automatically raises the MTU to 1554 bytes or 1574 bytes. In this case, the entire Ethernet datagram is being encapsulated, so the new packet is larger and requires a larger MTU. If the MTU used by other devices is larger, then you should set the source interface MTU to be the network MTU + 54 bytes for IPv4 or +64 bytes for IPv6. For the Firewall Threat Defense Virtual, this MTU requires a restart to enable jumbo frame reservation.
- Geneve encapsulation—If the source interface MTU is less than 1806 bytes, then the Firewall Threat Defense automatically raises the MTU to 1806 bytes. In this case, the entire Ethernet datagram is being encapsulated, so the new packet is larger and requires a larger MTU. If the MTU used by other devices is larger, then you should set the source interface MTU to be the network MTU + 306 bytes. This MTU requires a restart to enable jumbo frame reservation.

Configure VXLAN or Geneve Interfaces

You can configure either VXLAN or Geneve interfaces.

Configure VXLAN Interfaces

To configure VXLAN, perform the following steps.

Note You can configure either VXLAN or Geneve (Firewall Threat Defense Virtual only). For Geneve interfaces, see Configure Geneve Interfaces, on page 850.

Note For Azure GWLB, the VXLAN interface is configured when you deploy the VM using the ARM template. You can use this section to change your configuration.

1. Configure the VTEP Source Interface, on page 848.
2. Configure the VNI Interface, on page 849.
3. (Azure GWLB) Allow Gateway Load Balancer Health Checks, on page 852.

Configure the VTEP Source Interface

You can configure one VTEP source interface per Firewall Threat Defense device. The VTEP is defined as a Network Virtualization Endpoint (NVE). VXLAN is the default encapsulation type. An exception is made for clustering on the Firewall Threat Defense Virtual in Azure, where you can use one VTEP source interface for the cluster control link and a second one for the data interface connected to the Azure GWLB.

Procedure

Step 1 If you want to specify a group of peer VTEPs, add a network object with the peer IP addresses. See Creating Network Objects, on page 1403.

Step 2 Choose **Devices** > **Device Management**.

Step 3 Click **Edit** (✎) next to the device on which you want to configure VXLAN.

Step 4 (Optional) Specify that the source interface is NVE-only.

This setting is optional for routed mode where this setting restricts traffic to VXLAN and common management traffic only on this interface. This setting is automatically enabled for transparent firewall mode.

a) Click **Interfaces**.
b) Click **Edit** (✎) for the VTEP source interface.
c) On the **General** page, check the check box of **NVE Only**.

Step 5 Click **VTEP** if it is not already displaying.

Step 6 Check **Enable NVE**.

Step 7 Click **Add VTEP**.

Step 8 For the **Encapsulation Type**, choose **VxLAN**.

For AWS, you can choose between **VxLAN** and **Geneve**. Other platforms have **VxLAN** chosen automatically.

Step 9 Enter the value for the **Encapsulation port** within the specified range.

The default value is 4789.

Step 10 Select the **VTEP Source Interface**.

Select from the list of available physical interfaces present on the device. If the source interface MTU is less than 1554 bytes for IPv4 or 1574 bytes for IPv6, then the Firewall Management Center automatically raises the MTU to 1554 bytes or 1574 bytes.

Step 11 Select the **Neighbor Address**. The available options are:

- **None**—No neighbor address is specified.
- **Peer VTEP**—Specify a peer VTP address.
- **Peer Group**—Specify a network object with the peer IP addresses.
- **Default Multicast**—Specify a default multicast group for all associated VNI interfaces. If you do not configure the multicast group per VNI interface, then this group is used. If you configure a group at the VNI interface level, then that group overrides this setting.

Step 12 Click **OK**.

Step 13 Click **Save**.

Step 14 Configure the routed interface parameters. See Configure Routed Mode Interfaces.

Configure the VNI Interface

Add a VNI interface, associate it with the VTEP source interface, and configure basic interface parameters.

For the Firewall Threat Defense Virtual in Azure, you can configure either a regular VXLAN interface, or you can configure a paired proxy mode VXLAN interface for use with the Azure GWLB. Paired proxy mode is the only supported mode with clustering.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **Edit** (✎) next to the device on which you want to configure VXLAN.

Step 3 Click **Interfaces**.

Step 4 Click **Add Interfaces**, and then choose **VNI Interface**.

Step 5 Enter the interface **Name** and **Description**.

Step 6 From the **Security Zone** drop-down list, choose a security zone or add a new one by clicking **New**.

Step 7 Enter a value for the **Priority** field within the specified range. By default, 0 is selected.

Step 8 Enter a value for the **VNI ID** between 1 and 10000.

This ID is only an internal interface identifier.

Step 9 (Paired Proxy VXLAN for Azure GWLB) Enable proxy paired mode and set the required parameters.

a) Check **Proxy Paired**.
b) Set the **Internal Port** between 1024 and 65535.
c) Set the **Internal Segment ID** between 1 and 16777215.
d) Set the **External Port** between 1024 and 65535.
e) Set the **External Segment ID** between 1 and 16777215.

Step 10 (Regular VXLAN) Enter a value for the **VNI Segment ID** between 1 and 16777215.

The segment ID is used for VXLAN tagging.

Step 11 Enter the **Multicast Group IP Address**.

If you do not set the multicast group for the VNI interface, the default group from the VTEP source interface configuration is used, if available. If you manually set a VTEP peer IP for the VTEP source interface, you cannot specify a multicast group for the VNI interface.

Step 12 Check **NVE Mapped to VTEP Interface**.

This option associates this interface with the VTEP source interface.

Step 13 Click **OK**.

Step 14 Click **Save** to save the interface configuration.

Step 15 Configure the routed or transparent interface parameters. See Configure Routed and Transparent Mode Interfaces, on page 853.

Configure Geneve Interfaces

To configure Geneve interfaces for Firewall Threat Defense Virtual, perform the following steps.

Note You can configure either VXLAN or Geneve. For information about VXLAN interfaces, see Configure VXLAN Interfaces, on page 848.

1. Configure the VTEP Source Interface, on page 850.
2. Configure the VNI, on page 851.
3. Allow Gateway Load Balancer Health Checks, on page 852.

Configure the VTEP Source Interface

You can configure one VTEP source interface per Firewall Threat Defense Virtual device. The VTEP is defined as a Network Virtualization Endpoint (NVE).

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **Edit** (✎) next to the device on which you want to configure Geneve.

Step 3 Click **VTEP**.

Step 4 Check **Enable NVE**.

Step 5 Click **Add VTEP**.

Step 6 For the **Encapsulation Type**, choose **Geneve**.

Step 7 Enter the value for the **Encapsulation port** within the specified range.

We do not recommend changing the Geneve port; AWS requires a port of 6081.

Step 8 Select the **VTEP Source Interface**.

You can select from the list of available physical interfaces present on the device. If the source interface MTU is less than 1806 bytes, then the Firewall Management Center automatically raises the MTU to 1806 bytes.

Step 9 Click **OK**.

Step 10 Click **Save**.

Step 11 Configure the routed interface parameters. See Configure Routed Mode Interfaces.

Configure the VNI

Add a VNI, associate it with (VTEP) source interface, and configure basic interface parameters.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **Edit** (✎) next to the device on which you want to configure Geneve.

Step 3 Click **Interfaces**.

Step 4 Click **Add Interfaces**, and then choose **VNI Interface**.

Step 5 In the **Name** and **Description** fields, provide relevant information.

Step 6 In the **VNI ID** field, enter a value between 1 and 10000.

Note
This ID is only an internal interface identifier.

Step 7 Check the **Enable Proxy** check box.

Note
When the AWS proxy is enabled on the VNI interface of your device, configuring NAT is not allowed.

This option enables single-arm proxy mode or dual-arm proxy mode. Single-arm proxy mode allows traffic to exit the same interface it entered (U-turn traffic) while the dual-arm proxy mode enables the virtual devices to perform the NAT of the inspected traffic and then directly forward the outbound traffic to the internet without returning it to the GWLB and GWLB endpoint. If you edit the interface later, you cannot disable single-arm proxy mode or dual-arm proxy mode. To do that, you need to delete the existing interface and create a new VNI.

This option is only available for a Geneve VTEP.

Step 8 From the **Proxy type** drop-down list, choose the proxy mode you want to enable for the interface. If you do not specify the proxy mode for the interface, then by default, single-arm proxy mode is considered.

Step 9 Select **NVE Mapped to VTEP Interface**.

This option associates this interface with the VTEP source interface.

Step 10 Click **OK**.

Step 11 Click **Save**.

What to do next

Configure the routed interface parameters. See Configure Routed Mode Interfaces.

Allow Gateway Load Balancer Health Checks

The AWS or Azure GWLB requires appliances to answer a health check properly. The GWLB will only send traffic to appliances that are considered healthy. You must configure the Firewall Threat Defense Virtual to respond to an SSH, HTTP, or HTTPS health check.

Configure one of the following methods.

Procedure

Step 1 Configure SSH. See Configure Secure Shell

Allow SSH from the GWLB IP address. The GWLB will attempt to establish a connection to the Firewall Threat Defense Virtual, and the Firewall Threat Defense Virtual's prompt to log in is taken as proof of health. An SSH login attempt will time out after 1 minute. You will need to configure a longer health check interval on the GWLB to accommodate this timeout.

Step 2 Configure HTTP(S) Redirection Using Static Interface NAT with Port Translation.

You can configure the Firewall Threat Defense Virtual to redirect health checks to a metadata HTTP(S) server. For HTTP(S) health checks, the HTTP(S) server must reply to the GWLB with a status code in the range 200 to 399. Because the Firewall Threat Defense Virtual has limits on the number of simultaneous management connections, you may choose to offload the health check to an external server.

Static interface NAT with port translation lets you redirect a connection to a port (such as port 80) to a different IP address. For example, translate an HTTP packet from the GWLB with a destination of the Firewall Threat Defense Virtual outside interface so that it appears to be from the Firewall Threat Defense Virtual outside interface with a destination of the HTTP server. The Firewall Threat Defense Virtual then forwards the packet to the mapped destination address. The HTTP server responds to the Firewall Threat Defense Virtual outside interface, and then the Firewall Threat Defense Virtual forwards the response back to the GWLB. You need an access rule that allows traffic from the GWLB to the HTTP server.

a) Permit HTTP(S) traffic on the outside interface from the GWLB network in an access rule. See Access Control Rules, on page 1773.

b) For HTTP(S), translate the source GWLB IP address to the Firewall Threat Defense Virtual outside interface IP address; then translate the destination of the outside interface IP address to the HTTP(S) server IP address. See Configure Static Manual NAT, on page 1058.

Configure Routed and Transparent Mode Interfaces

This section includes tasks to complete the regular interface configuration for all models in routed or transparent firewall mode.

About Routed and Transparent Mode Interfaces

Firewall mode interfaces subject traffic to firewall functions such as maintaining flows, tracking flow states at both IP and TCP layers, IP defragmentation, and TCP normalization. You can also optionally configure IPS functions for this traffic according to your security policy.

The types of firewall interfaces you can configure depends on the firewall mode set for the device: routed or transparent mode. See Transparent or Routed Firewall Mode, on page 285 for more information.

- Routed mode interfaces (routed firewall mode only)—Each interface that you want to route between is on a different subnet.
- Bridge group interfaces (routed and transparent firewall mode)—You can group together multiple interfaces on a network, and the Firewall Threat Defense device uses bridging techniques to pass traffic between the interfaces. Each bridge group includes a Bridge Virtual Interface (BVI) to which you assign an IP address on the network. In routed mode, the Firewall Threat Defense device routes between BVIs and regular routed interfaces. In transparent mode, each bridge group is separate and cannot communicate with each other.

Dual IP Stack (IPv4 and IPv6)

The Firewall Threat Defense device supports both IPv6 and IPv4 addresses on an interface. Make sure you configure a default route for both IPv4 and IPv6.

31-Bit Subnet Mask

For routed interfaces, you can configure an IP address on a 31-bit subnet for point-to-point connections. The 31-bit subnet includes only 2 addresses; normally, the first and last address in the subnet is reserved for the network and broadcast, so a 2-address subnet is not usable. However, if you have a point-to-point connection and do not need network or broadcast addresses, a 31-bit subnet is a useful way to preserve addresses in IPv4. For example, the failover link between 2 Firewall Threat Defenses only requires 2 addresses; any packet that is transmitted by one end of the link is always received by the other, and broadcasting is unnecessary. You can also have a directly-connected management station running SNMP or Syslog.

31-Bit Subnet and Clustering

You can use a 31-bit subnet mask for cluster interfaces, excluding the management interface and the Cluster Control Link.

31-Bit Subnet and Failover

For failover, when you use a 31-bit subnet for the Firewall Threat Defense interface IP address, you cannot configure a standby IP address for the interface because there are not enough addresses. Normally, an interface for failover should have a standby IP address so the active unit can perform interface tests to ensure standby interface health. Without a standby IP address, the Firewall Threat Defense cannot perform any network tests; only the link state can be tracked.

For the failover and optional separate state link, which are point-to-point connections, you can also use a 31-bit subnet.

31-Bit Subnet and Management

If you have a directly-connected management station, you can use a point-to-point connection for SSH or HTTP on the Firewall Threat Defense, or for SNMP or Syslog on the management station.

31-Bit Subnet Unsupported Features

The following features do not support the 31-Bit subnet:

- BVI interfaces for bridge groups—The bridge group requires at least 3 host addresses: the BVI, and two hosts connected to two bridge group member interfaces. you must use a /29 subnet or smaller.
- Multicast Routing

Guidelines and Limitations for Routed and Transparent Mode Interfaces

High Availability, Clustering, and Multi-Instance

- Do not configure failover links with the procedures in this chapter. See the High Availability chapter for more information.
- For cluster interfaces, see the clustering chapter for requirements.
- For multi-instance mode, shared interfaces are not supported for bridge group member interfaces (in transparent mode or routed mode).
- When you use High Availability, you must set the IP address and standby address for data interfaces manually; DHCP and PPPoE are not supported. Set the standby IP addresses on the **Devices** > **Device Management** > **High Availability** tab in the **Monitored Interfaces** area. See the High Availability chapter for more information.

IPv6

- IPv6 is supported on all interfaces.
- You can only configure IPv6 addresses manually in transparent mode.
- The Firewall Threat Defense device does not support IPv6 anycast addresses.
- DHCPv6 and prefix delegation options are not supported with transparent mode, clustering, or High Availability.

Model Guidelines

- For the Firewall Threat Defense Virtual on VMware with bridged ixgbevf interfaces, bridge groups are not supported.

Transparent Mode and Bridge Group Guidelines

- You can create up to 250 bridge groups, with 64 interfaces per bridge group.

- Each directly-connected network must be on the same subnet.
- The Firewall Threat Defense device does not support traffic on secondary networks; only traffic on the same network as the BVI IP address is supported.
- An IP address for the BVI is required for each bridge group for to-the-device and from-the-device management traffic, as well as for data traffic to pass through the Firewall Threat Defense device. For IPv4 traffic, specify an IPv4 address. For IPv6 traffic, specify an IPv6 address.
- You can only configure IPv6 addresses manually.
- The BVI IP address must be on the same subnet as the connected network. You cannot set the subnet to a host subnet (255.255.255.255).
- Management interfaces are not supported as bridge group members.
- For multi-instance mode, shared interfaces are not supported for bridge group member interfaces (in transparent mode or routed mode).
- For the Firewall Threat Defense Virtual on VMware with bridged ixgbevf interfaces, transparent mode is not supported, and bridge groups are not supported in routed mode.
- For the Firepower 1010 and Secure Firewall 1210/20, you cannot mix logical VLAN interfaces and physical firewall interfaces in the same bridge group.
- For the Firepower 4100/9300, data-sharing interfaces are not supported as bridge group members.
- In transparent mode, you must use at least 1 bridge group; data interfaces must belong to a bridge group.
- In transparent mode, do not specify the BVI IP address as the default gateway for connected devices; devices need to specify the router on the other side of the Firewall Threat Defense as the default gateway.
- In transparent mode, the *default* route, which is required to provide a return path for management traffic, is only applied to management traffic from one bridge group network. This is because the default route specifies an interface in the bridge group as well as the router IP address on the bridge group network, and you can only define one default route. If you have management traffic from more than one bridge group network, you need to specify a regular static route that identifies the network from which you expect management traffic.
- Transparent mode is not supported on threat defense virtual instances deployed on Amazon Web Services, Microsoft Azure, Google Cloud Platform, and Oracle Cloud Infrastructure.
- In routed mode, to route between bridge groups and other routed interfaces, you must name the BVI.
- In routed mode, Firewall Threat Defense-defined EtherChannel interfaces are not supported as bridge group members. EtherChannels on the Firepower 4100/9300 can be bridge group members.
- Bidirectional Forwarding Detection (BFD) echo packets are not allowed through the Firewall Threat Defense when using bridge group members. If there are two neighbors on either side of the Firewall Threat Defense running BFD, then the Firewall Threat Defense will drop BFD echo packets because they have the same source and destination IP address and appear to be part of a LAND attack.

Additional Guidelines and Requirements

- The Firewall Threat Defense supports only one 802.1Q header in a packet and does not support multiple headers (known as Q-in-Q support) for firewall interfaces.**Note:** For inline sets and passive interfaces,

the FTD supports Q-in-Q up to two 802.1Q headers in a packet, with the exception of the Firepower 4100/9300, which only supports one 802.1Q header.

- Interface problems, such as frequent up/down status changes, can prevent the floating connection timer from applying correctly to the connections going through the interface. If you have problems with an interface's status, consider clearing all connections after the status becomes stable to clear invalid connections.

Configure Routed Mode Interfaces

This procedure describes how to set the name, security zone, and IPv4 address.

Note Not all fields are supported for all interface types.

Before you begin

- **Firepower 4100/9300**
  1. Configure a Physical Interface, on page 332
  2. (Optional) Configure any special interfaces.
     - Add an EtherChannel (Port Channel), on page 333
     - Add a VLAN Subinterface for Container Instances, on page 335 in FXOS
     - Configure a Loopback Interface, on page 830
     - Add a Subinterface, on page 836 in Firewall Management Center
     - Configure VXLAN Interfaces, on page 848

- (Optional) **All other models:**
  - Configure an EtherChannel, on page 786
  - Configure a Loopback Interface, on page 830
  - Add a Subinterface, on page 836
  - Configure VXLAN Interfaces, on page 848
  - Firewall Threat Defense Virtual on AWS: Configure Geneve Interfaces, on page 850
  - Firepower 1010 and Secure Firewall 1210/1220: Configure a VLAN Interface, on page 821

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 In the **Name** field, enter a name up to 48 characters in length.

You cannot start the name with the phrase "cluster". It is reserved for internal use.

Step 4 Enable the interface by checking the **Enabled** check box.

Step 5 (Optional) Set this interface to **Management Only** to limit traffic to management traffic; through-the-box traffic is not allowed.

Step 6 (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

Step 7 In the **Mode** drop-down list, choose **None**.

Regular firewall interfaces have the mode set to None. The other modes are for IPS-only interface types.

Step 8 From the **Security Zone** drop-down list, choose a security zone or add a new one by clicking **New**.

The routed interface is a Routed-type interface, and can only belong to Routed-type zones.

Step 9 See Configure the MTU, on page 880 for information about the **MTU**.

Step 10 In the **Priority** field, enter a number ranging from 0–65535.

This value is used in the policy based routing configuration. The priority is used to determine how you want to route the traffic across multiple egress interfaces. For more information, see Configure Policy-Based Routing Policy, on page 1329.

Step 11 Click the **IPv4** tab. To set the IP address, use one of the following options from the **IP Type** drop-down list.

High Availability, clustering, and loopback interfaces only support static IP address configuration; DHCP and PPPoE are not supported.

- **Use Static IP**—Enter the IP address and subnet mask. For point-to-point connections, you can specify a 31-bit subnet mask (255.255.255.254 or /31). In this case, no IP addresses are reserved for the network or broadcast addresses. You cannot set the standby IP address in this case. For High Availability, you can only use a static IP address. Set the standby IP address on the **Devices** > **Device Management** page, click **High Availability** in the **Monitored Interfaces** area. If you do not set the standby IP address, the active unit cannot monitor the standby interface using network tests; it can only track the link state.
- **Use DHCP**—Configure the following optional parameters:
  - **Obtain default route using DHCP**—Obtains the default route from the DHCP server.
  - **DHCP route metric**—Assigns an administrative distance to the learned route, between 1 and 255. The default administrative distance for the learned routes is 1.
- **Use PPPoE**—If the interface is connected to a DSL, cable modem, or other connection to your ISP, and your ISP uses PPPoE to provide your IP address, configure the following parameters:
  - **VPDN Group Name**—Specify a group name of your choice to represent this connection.
  - **PPPoE User Name**—Specify the username provided by your ISP.
  - **PPPoE Password/Confirm Password**—Specify and confirm the password provided by your ISP.
  - **PPP Authentication**—Choose **PAP**, **CHAP**, or **MSCHAP**.

PAP passes a cleartext username and password during authentication and is not secure. With CHAP, the client returns the encrypted [challenge plus password], with a cleartext username in response to the server challenge. CHAP is more secure than PAP, but it does not encrypt data. MSCHAP is similar to CHAP but is more secure because the server stores and compares only encrypted passwords rather than cleartext passwords as in CHAP. MSCHAP also generates a key for data encryption by MPPE.

- **PPPoE route metric**—Assign an administrative distance to the learned route. Valid values are from 1 to 255. By default, the administrative distance for the learned routes is 1.
- **Enable Route Settings**—To manually configure the PPPoE IP address, check this box and then enter the **IP Address**.

  If you select the **Enable Route Settings** check box and leave the **IP Address** blank, the **ip address pppoe setroute** command is applied as shown in this example:

```
interface GigabitEthernet0/2
nameif inside2_pppoe
cts manual
  propagate sgt preserve-untag
  policy static sgt disabled trusted
security-level 0
pppoe client vpdn group test
pppoe client route distance 10
ip address pppoe setroute
```

- **Store Username and Password in Flash**—Stores the username and password in flash memory.

  The Firewall Threat Defense device stores the username and password in a special location of NVRAM.

Step 12 (Optional) See Configure IPv6 Addressing, on page 864 to configure IPv6 addressing on the **IPv6** tab.

Step 13 (Optional) See Configure the MAC Address, on page 881 to manually configure the MAC address on the **Advanced** tab.

Step 14 (Optional) Set the duplex and speed by clicking **Hardware Configuration** > **Speed**.

- **Duplex**—Choose **Full** or **Half**. SFP interfaces only support **Full** duplex.
- **Speed**—Choose a speed (varies depending on the model). (Secure Firewall 3100/4200 only) Choose **Detect SFP** to detect the speed of the installed SFP module and use the appropriate speed. Duplex is always Full, and auto-negotiation is always enabled. This option is useful if you later change the network module to a different model, and want the speed to update automatically.For Secure Firewall 1250, you can configure a maximum interface speed of 2.5gbps.

  Note
  You cannot modify the speed of a HA or a Cluster Control Link interface.

- **Auto-negotiation**—Set the interface to negotiate the speed, link status, and flow control.
- **Forward Error Correction Mode**—(Secure Firewall 3100/4200 only) For 25 Gbps and higher interfaces, enable Forward Error Correction (FEC). For an EtherChannel member interface, you must configure FEC before you add it to the EtherChannel. The setting chosen when you use **Auto** depends on the transceiver type and whether the interface is fixed (built-in) or on a network module.

Table 52: Default FEC for Auto Setting

Transceiver Type	Fixed Port Default FEC (Ethernet 1/9 through 1/16)	Network Module Default FEC
25G-SR	Clause 108 RS-FEC	Clause 108 RS-FEC
25G-LR	Clause 108 RS-FEC	Clause 108 RS-FEC
10/25G-CSR	Clause 108 RS-FEC	Clause 74 FC-FEC
25G-AOC*x*M	Clause 74 FC-FEC	Clause 74 FC-FEC
25G-CU2.5/3M	Auto-Negotiate	Auto-Negotiate
25G-CU4/5M	Auto-Negotiate	Auto-Negotiate
25/50/100G	Clause 91 RS-FEC	Clause 91 RS-FEC

Step 15 (Optional) Enable Firewall Management Center manager access on a data interface on the **Manager Access** page.

You can enable manager access from a data interface when you first setup the Firewall Threat Defense. If you want to enable or disable manager access after you added the Firewall Threat Defense to the Firewall Management Center, see:

- Enable manager access: Change the Manager Access Interface from Management to Data, on page 177

  Note
  You cannot enable manager access unless you first initiate the manager interface migration from Management to a data interface. After you initiate the migration, you can enable manager access on the **Manager Access** page and save the configuration successfully.

- Disable manager access: Change the Manager Access Interface from Data to Management, on page 182

If you want to change the manager access interface from one data interface to another data interface, you must disable manager access on the original data interface, but do not disable the interface itself yet; the original data interface must be used to perform the deployment. If you want to use the same IP address on the new manager access interface, you can delete or change the IP configuration on the original interface; this change should not affect the deployment. If you use a different IP address for the new interface, then also change the device IP address shown in the Firewall Management Center; see Update the Hostname or IP Address in the Firewall Management Center, on page 170. Be sure to also update related configuration to use the new interface such as static routes, DDNS, and DNS settings.

Manager access from a data interface has the following limitations:

- You can only enable manager access on a physical, data interface. You cannot use a subinterface or EtherChannel, nor can you create a subinterface on the manager access interface. You can also use the Firewall Management Center to enable manager access on a single secondary interface for redundancy.
- This interface cannot be management-only.
- Routed firewall mode only, using a routed interface.
- PPPoE is not supported. If your ISP requires PPPoE, you will have to put a router with PPPoE support between the Firewall Threat Defense and the WAN modem.

- The interface must be in the global VRF only.
- SSH is not enabled by default for data interfaces, so you will have to enable SSH later using the Firewall Management Center. Because the Management interface gateway will be changed to be the data interfaces, you also cannot SSH to the Management interface from a remote network unless you add a static route for the Management interface using the **configure network static-routes** command. For Firewall Threat Defense Virtual on Amazon Web Services, a console port is not available, so you should maintain your SSH access to the Management interface: add a static route for Management before you continue with your configuration. Alternatively, be sure to finish all CLI configuration (including the **configure manager add** command) before you configure the data interface for manager access and you are disconnected.
- You cannot use separate management and event-only interfaces.
- Clustering is not supported. You must use the Management interface in this case.

Figure 377: Manager Access

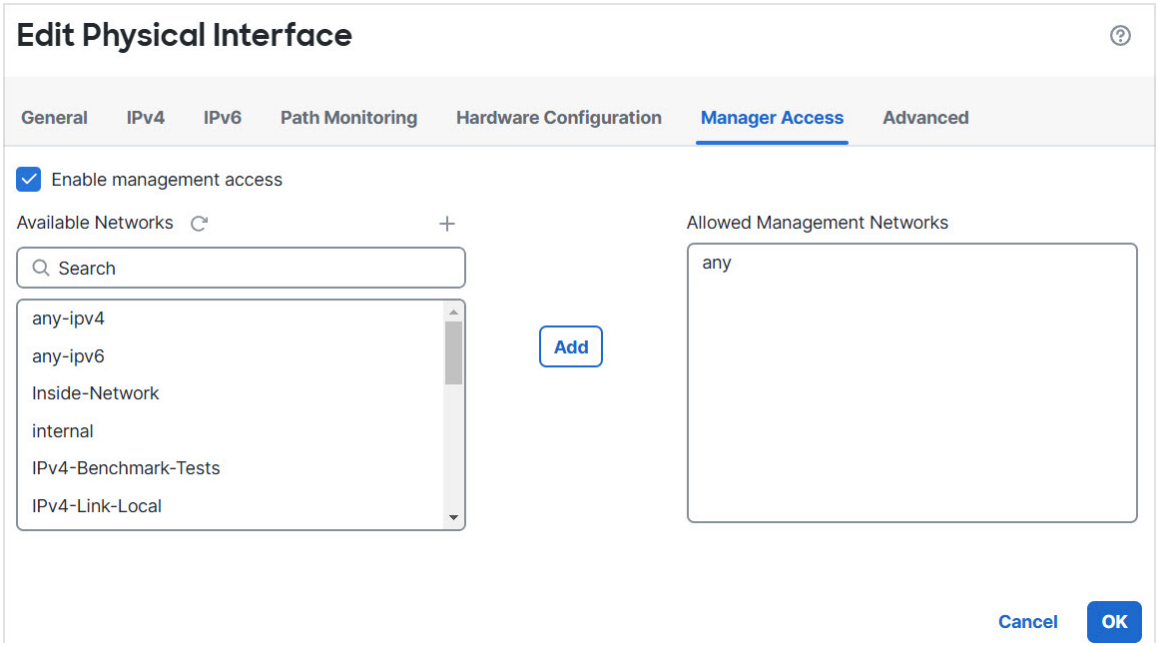

- Check **Enable management on this interface for the manager** to use this data interface for management instead of the dedicated Management interface.
- (Optional) In the **Allowed Management Networks** box, add the networks from which you want to allow manager access. By default, any networks are allowed.

Step 16 Click **OK**.

Step 17 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Bridge Group Interfaces

A bridge group is a group of interfaces that the Secure Firewall Threat Defense device bridges instead of routes. Bridge groups are supported in both transparent and routed firewall mode. For more information about bridge groups, see About Bridge Groups, on page 287.

To configure bridge groups and associated interfaces, perform these steps.

Configure General Bridge Group Member Interface Parameters

This procedure describes how to set the name and security zone for each bridge group member interface. The same bridge group can include different types of interfaces: physical interfaces, VLAN subinterfaces, Firepower 1010 and Secure Firewall 1210/1220 VLAN interfaces, EtherChannels, and redundant interfaces. The Management interface is not supported. In routed mode, EtherChannels are not supported. For the Firepower 4100/9300, data-sharing type interfaces are not supported.

Before you begin

- **Firepower 4100/9300**
  1. Configure a Physical Interface, on page 332
  2. (Optional) Configure any special interfaces.
     - Add an EtherChannel (Port Channel), on page 333
     - Add a VLAN Subinterface for Container Instances, on page 335 in FXOS
     - Add a Subinterface, on page 836 in Firewall Management Center
- (Optional) **All other models:**
  - Configure an EtherChannel, on page 786
  - Add a Subinterface, on page 836
  - Firepower 1010 and Secure Firewall 1210/1220: Configure a VLAN Interface, on page 821

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 In the **Name** field, enter a name up to 48 characters in length.

You cannot start the name with the phrase "cluster". It is reserved for internal use.

Step 4 Enable the interface by checking the **Enabled** check box.

Step 5 (Optional) Set this interface to **Management Only** to limit traffic to management traffic; through-the-box traffic is not allowed.

Step 6 (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

Step 7 In the **Mode** drop-down list, choose **None**.

Regular firewall interfaces have the mode set to None. The other modes are for IPS-only interface types. After you assign this interface to a bridge group, the mode will show as **Switched**.

Step 8 From the **Security Zone** drop-down list, choose a security zone or add a new one by clicking **New**.

The bridge group member interface is a Switched-type interface, and can only belong to Switched-type zones. Do not configure any IP address settings for this interface. You will set the IP address for the Bridge Virtual Interface (BVI) only. Note that the BVI does not belong to a zone, and you cannot apply access control policies to the BVI.

Step 9 See Configure the MTU, on page 880 for information about the **MTU**.

Step 10 (Optional) Set the duplex and speed by clicking **Hardware Configuration** > **Speed**.

- **Duplex**—Choose **Full** or **Half**. SFP interfaces only support **Full** duplex.
- **Speed**—Choose a speed (varies depending on the model). (Secure Firewall 3100/4200 only) Choose **Detect SFP** to detect the speed of the installed SFP module and use the appropriate speed. Duplex is always Full, and auto-negotiation is always enabled. This option is useful if you later change the network module to a different model, and want the speed to update automatically.For Secure Firewall 1250, you can configure a maximum interface speed of 2.5gbps.

  Note
  You cannot modify the speed of a HA or a Cluster Control Link interface.

- **Auto-negotiation**—Set the interface to negotiate the speed, link status, and flow control.
- **Forward Error Correction Mode**—(Secure Firewall 3100/4200 only) For 25 Gbps and higher interfaces, enable Forward Error Correction (FEC). For an EtherChannel member interface, you must configure FEC before you add it to the EtherChannel. The setting chosen when you use **Auto** depends on the transceiver type and whether the interface is fixed (built-in) or on a network module.

Table 53: Default FEC for Auto Setting

Transceiver Type	Fixed Port Default FEC (Ethernet 1/9 through 1/16)	Network Module Default FEC
25G-SR	Clause 108 RS-FEC	Clause 108 RS-FEC
25G-LR	Clause 108 RS-FEC	Clause 108 RS-FEC
10/25G-CSR	Clause 108 RS-FEC	Clause 74 FC-FEC
25G-AOC*x*M	Clause 74 FC-FEC	Clause 74 FC-FEC
25G-CU2.5/3M	Auto-Negotiate	Auto-Negotiate
25G-CU4/5M	Auto-Negotiate	Auto-Negotiate
25/50/100G	Clause 91 RS-FEC	Clause 91 RS-FEC

Step 11 (Optional) See Configure IPv6 Addressing, on page 864 to configure IPv6 addressing on the **IPv6** tab.

Step 12 (Optional) See Configure the MAC Address, on page 881 to manually configure the MAC address on the **Advanced** tab.

Step 13 Click **OK**.

Step 14 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure the Bridge Virtual Interface (BVI)

Each bridge group requires a BVI for which you configure an IP address. The Firewall Threat Defense uses this IP address as the source address for packets originating from the bridge group. The BVI IP address must be on the same subnet as the connected network. For IPv4 traffic, the BVI IP address is required to pass any traffic. For IPv6 traffic, you must, at a minimum, configure the link-local addresses to pass traffic, but a global management address is recommended for full functionality, including remote management and other management operations.

For routed mode, if you provide a name for the BVI, then the BVI participates in routing. Without a name, the bridge group remains isolated as in transparent firewall mode.

Before you begin

You cannot add the BVI to a security zone; therefore, you cannot apply Access Control policies to the BVI. You must apply your policy to the bridge group member interfaces based on their zones.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Choose **Add Interfaces** > **Bridge Group Interface**.

Step 3 (Routed Mode) In the **Name** field, enter a name up to 48 characters in length.

You must name the BVI if you want to route traffic outside the bridge group members, for example, to the outside interface or to members of other bridge groups. The name is not case-sensitive.

Step 4 In the **Bridge Group ID** field, enter the bridge group ID between 1 and 250.

Step 5 In the **Description** field, enter a description for this bridge group.

Step 6 On the **Interfaces** tab, click an interface and then click **Add** to move it to the **Selected Interfaces** area. Repeat for all interfaces that you want to make members of the bridge group.

Step 7 (Transparent Mode) Click the **IPv4** tab. In the **IP Address** field, enter the IPv4 address and subnet mask.

Do not assign a host address (/32 or 255.255.255.255) to the BVI. Also, do not use other subnets that contain fewer than 3 host addresses (one each for the upstream router, downstream router, and transparent firewall) such as a /30 subnet (255.255.255.252). The Firewall Threat Defense device drops all ARP packets to or from the first and last addresses in a subnet. For example, if you use a /30 subnet and assign a reserved address from that subnet to the upstream router, then the Firewall Threat Defense device drops the ARP request from the downstream router to the upstream router.

For High Availability, set the standby IP address on the **Devices** > **Device Management** page, click the **High Availability** tab in the **Monitored Interfaces** area. If you do not set the standby IP address, the active unit cannot monitor the standby interface using network tests; it can only track the link state.

Step 8 (Routed Mode) Click the **IPv4** tab. To set the IP address, use one of the following options from the **IP Type** drop-down list.

High Availability and clustering interfaces only support static IP address configuration; DHCP is not supported.

- **Use Static IP**—Enter the IP address and subnet mask. For High Availability, you can only use a static IP address. Set the standby IP address on the **Devices** > **Device Management**, click the **High Availability** tab in the **Monitored Interfaces** area. If you do not set the standby IP address, the active unit cannot monitor the standby interface using network tests; it can only track the link state.
- **Use DHCP**—Configure the following optional parameters:
  - **Obtain default route using DHCP**—Obtains the default route from the DHCP server.
  - **DHCP route metric**—Assigns an administrative distance to the learned route, between 1 and 255. The default administrative distance for the learned routes is 1.

Step 9 (Optional) See Configure IPv6 Addressing, on page 864 to configure IPv6 addressing.

Step 10 (Optional) See Add a Static ARP Entry, on page 882 and Add a Static MAC Address and Disable MAC Learning for a Bridge Group, on page 883 (for transparent mode only) to configure the **ARP** and **MAC** settings.

Step 11 Click **OK**.

Step 12 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure IPv6 Addressing

This section describes how to configure IPv6 addressing in routed and transparent mode.

About IPv6

This section includes information about IPv6.

IPv6 Addressing

You can configure two types of unicast addresses for IPv6:

- Global—The global address is a public address that you can use on the public network. For a bridge group, this address needs to be configured for the BVI, and not per member interface. You can also configure a global IPv6 address for the management interface in transparent mode.
- Link-local—The link-local address is a private address that you can only use on the directly-connected network. Routers do not forward packets using link-local addresses; they are only for communication on a particular physical network segment. They can be used for address configuration or for the Neighbor Discovery functions such as address resolution. In a bridge group, only member interfaces have link-local addresses; the BVI does not have a link-local address.

At a minimum, you need to configure a link-local address for IPv6 to operate. If you configure a global address, a link-local address is automatically configured on the interface, so you do not also need to specifically configure a link-local address. For bridge group member interfaces, when you configure the global address on the BVI, the Firewall Threat Defense device automatically generates link-local addresses for member interfaces. If you do not configure a global address, then you need to configure the link-local address, either automatically or manually.

Modified EUI-64 Interface IDs

RFC 3513: Internet Protocol Version 6 (IPv6) Addressing Architecture requires that the interface identifier portion of all unicast IPv6 addresses, except those that start with binary value 000, be 64 bits long and be constructed in Modified EUI-64 format. The Firewall Threat Defense device can enforce this requirement for hosts attached to the local link.

When this feature is enabled on an interface, the source addresses of IPv6 packets received on that interface are verified against the source MAC addresses to ensure that the interface identifiers use the Modified EUI-64 format. If the IPv6 packets do not use the Modified EUI-64 format for the interface identifier, the packets are dropped and the following system log message is generated:

```
325003: EUI-64 source address check failed.
```

The address format verification is only performed when a flow is created. Packets from an existing flow are not checked. Additionally, the address verification can only be performed for hosts on the local link.

Configure the IPv6 Prefix Delegation Client

The Firewall Threat Defense can act as a DHPCv6 Prefix Delegation client so that the client interface, for example the outside interface connected to a cable modem, can receive one or more IPv6 prefixes that the Firewall Threat Defense can then subnet and assign to its inside interfaces.

About IPv6 Prefix Delegation

The Firewall Threat Defense can act as a DHPCv6 Prefix Delegation client so that the client interface, for example the outside interface connected to a cable modem, can receive one or more IPv6 prefixes that the Firewall Threat Defense can then subnet and assign to its inside interfaces. Hosts connected to the inside interfaces can then use StateLess Address Auto Configuration (SLAAC) to obtain global IPv6 addresses. Note that the inside Firewall Threat Defense interfaces do not in turn act as Prefix Delegation servers; the Firewall Threat Defense can only provide global IP addresses to SLAAC clients. For example, if a router is connected to the Firewall Threat Defense, it can act as a SLAAC client to obtain its IP address. But if you want to use a subnet of the delegated prefix for the networks behind the router, you must manually configure those addresses on the router's inside interfaces.

The Firewall Threat Defense includes a light DHCPv6 server so the Firewall Threat Defense can provide information such as the DNS server and domain name to SLAAC clients when they send Information Request (IR) packets to the Firewall Threat Defense. The Firewall Threat Defense only accepts IR packets, and does not assign addresses to the clients. You will configure the client to generate its own IPv6 address by enabling IPv6 autoconfiguration on the client. Enabling stateless autoconfiguration on a client configures IPv6 addresses based on prefixes received in Router Advertisement messages; in other words, based on the prefix that the Firewall Threat Defense received using Prefix Delegation.

IPv6 Prefix Delegation /64 Subnet Example

The following example shows the Firewall Threat Defense receiving an IP address on the outside interface using the DHCPv6 address client. It also gets a delegated prefix using the DHCPv6 Prefix Delegation client.

The Firewall Threat Defense subnets the delegated prefix into /64 networks and assigns global IPv6 addresses to its inside interfaces dynamically using the delegated prefix plus a manually configured subnet (::0, ::1, or ::2) and IPv6 address (0:0:0:1) per interface. SLAAC clients connected to those inside interfaces obtain IPv6 addresses on each /64 subnet.

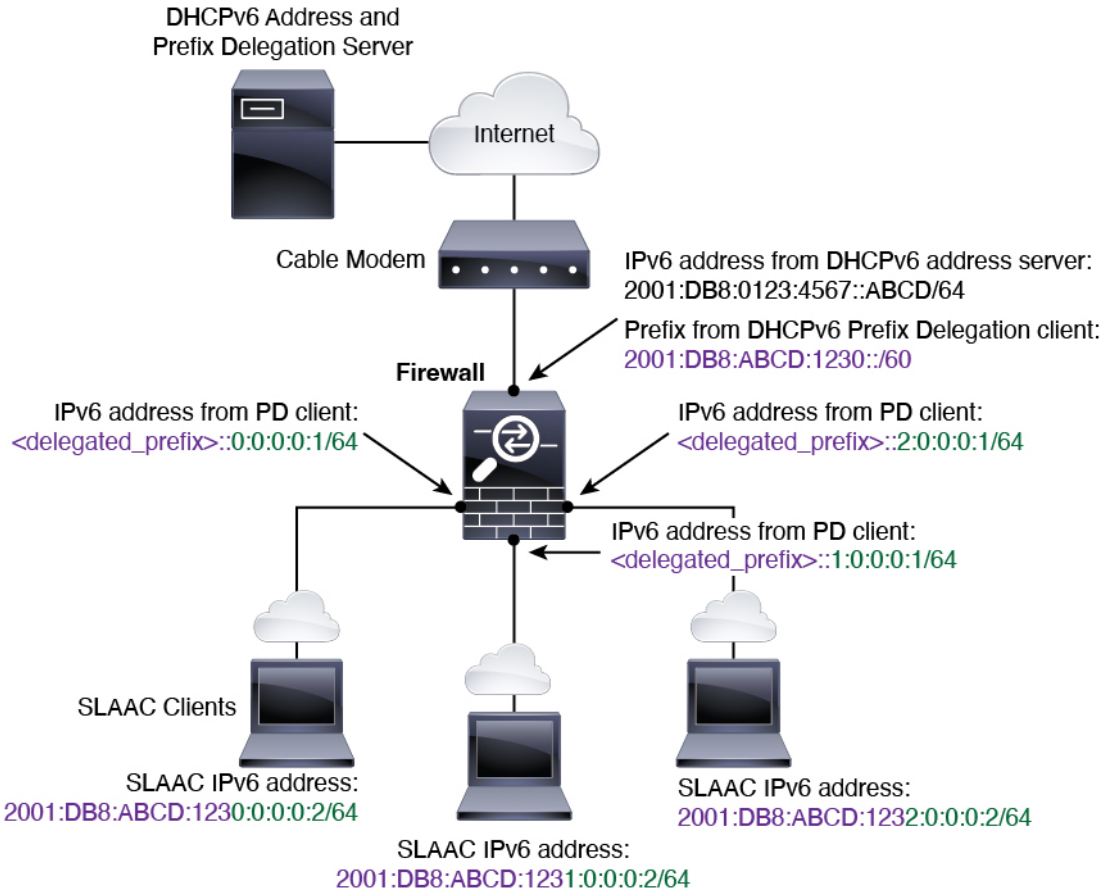

IPv6 Prefix Delegation /62 Subnet Example

The following example shows the Firewall Threat Defense subnetting the prefix into 4 /62 subnets: 2001:DB8:ABCD:123**0**::/62, 2001:DB8:ABCD:123**4**::/62, 2001:DB8:ABCD:123**8**::/62, and 2001:DB8:ABCD:123**C**::/62. The Firewall Threat Defense uses one of 4 available /64 subnets on 2001:DB8:ABCD:1230::/62 for its inside network (::0). You can then manually use additional /62 subnets for downstream routers. The router shown uses 3 of 4 available /64 subnets on 2001:DB8:ABCD:1234::/62 for its inside interfaces (::4, ::5, and ::6). In this case, the inside router interfaces cannot dynamically obtain the delegated prefix, so you need to view the delegated prefix on the Firewall Threat Defense, and then use that prefix for your router configuration. Usually, ISPs delegate the same prefix to a given client when the lease expires, but if the Firewall Threat Defense receives a new prefix, you will have to modify the router configuration to use the new prefix. The DHCP unique identifier (DUID) is persistent across reboots.

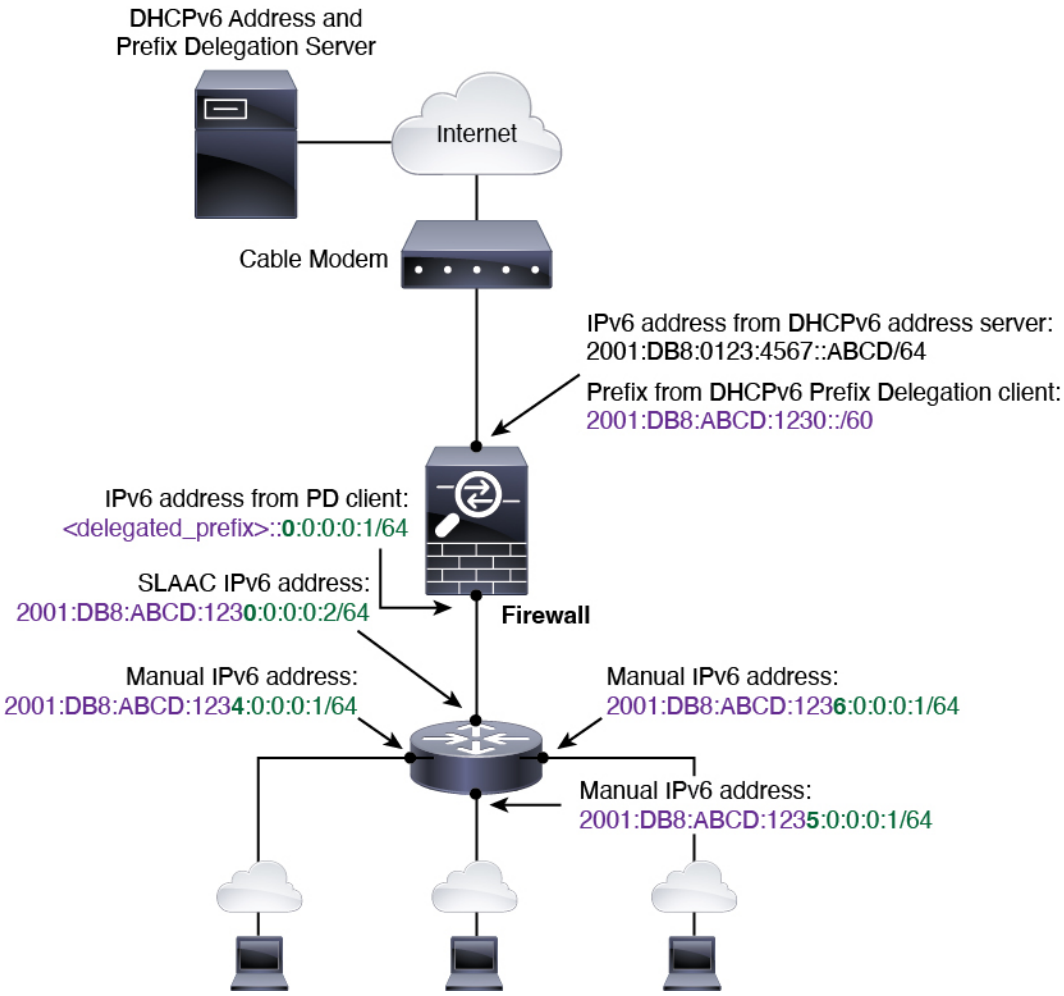

Enable the IPv6 Prefix Delegation Client

Enable the DHCPv6 Prefix Delegation client on one or more interfaces. The Firewall Threat Defense obtains one or more IPv6 prefixes that it can subnet and assign to inside networks. Typically, the interface on which you enable the prefix delegation client obtains its IP address using the DHCPv6 address client; only other Firewall Threat Defense interfaces use addresses derived from the delegated prefix.

This feature is only supported in routed mode. This feature is not supported in clustering or High Availability.

Before you begin

When you use Prefix Delegation, you must set the Firewall Threat Defense IPv6 neighbor discovery router advertisement interval to be much lower than the preferred lifetime of the prefix assigned by the DHCPv6 Server to prevent IPv6 traffic interruption. For example, if the DHCPv6 server sets the preferred Prefix Delegation lifetime to 300 seconds, you should set the Firewall Threat Defense RA interval to be 150 seconds. To set the preferred lifetime, use the **show ipv6 general-prefix** command. To set the Firewall Threat Defense RA interval, see Configure IPv6 Neighbor Discovery, on page 873; the default is 200 seconds.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click the **IPv6** page, and then click **DHCP**.

Step 4 Click **Client PD Prefix Name** and enter a name for this prefix.

Figure 378: Enable the Prefix Delegation Client

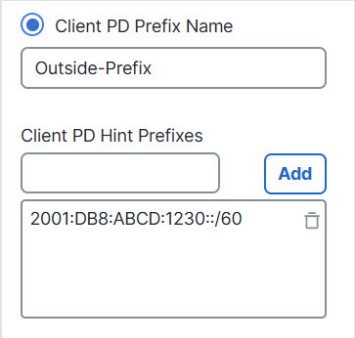

The name can be up to 200 characters.

Step 5 (Optional) Enter the prefix and prefix length in the **Client PD Hint Prefixes** field to provide one or more hints to the DHCP server about the delegated prefix you want to receive, then click **Add**.

Typically you want to request a particular prefix length, such as ::/60, or if you have received a particular prefix before and want to ensure you get it again when the lease expires, you can enter the whole prefix as the hint. If you enter multiple hints (different prefixes or lengths), then it is up to the DHCP server which hint to honor, or whether to honor the hint at all.

Step 6 Click **OK**.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure a Global IPv6 Address

To configure a global IPv6 address for any routed mode interface and for the transparent or routed mode BVI, perform the following steps.

Note Configuring the global address automatically configures the link-local address, so you do not need to configure it separately. For bridge groups, configuring the global address on the BVI automatically configures link-local addresses on all member interfaces.

For subinterfaces defined on the Firewall Threat Defense, we recommend that you also set the MAC address manually, because they use the same burned-in MAC address of the parent interface. IPv6 link-local addresses are generated based on the MAC address, so assigning unique MAC addresses to subinterfaces allows for unique IPv6 link-local addresses, which can avoid traffic disruption in certain instances on the Firewall Threat Defense. See Configure the MAC Address, on page 881.

Before you begin

For IPv6 neighbor discovery for bridge groups, you must explicitly allow Neighbor Solicitation (ICMPv6 type 135) and Neighbor Advertisement (ICMPv6 type 136) packets through the Firewall Threat Defense bridge group member interfaces using a bidirectional access rule.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click the **IPv6** page.

For routed mode, the **Basic** page is selected by default. For transparent mode, the **Address** page is selected by default.

Step 4 (Optional) On the **Basic** page, check **Enable IPv6**.

Use this option if you want to only configure the link-local addresses. Otherwise, configuring an IPv6 address enabled IPv6 processing automatically.

Step 5 Configure the global IPv6 address using one of the following methods.

For failover and clustering, and for loopback interfaces, you must set the IP address manually. For clustering, manually configuring the link-local address is also not supported.

- (Routed interface) Stateless autoconfiguration—Check the **Autoconfiguration** check box.

  Enabling stateless autoconfiguration on the interface configures IPv6 addresses based upon prefixes received in Router Advertisement messages. A link-local address, based on the Modified EUI-64 interface ID, is automatically generated for the interface when stateless autoconfiguration is enabled.

  Although RFC 4862 specifies that hosts configured for stateless autoconfiguration do not send Router Advertisement messages, the Firewall Threat Defense device does send Router Advertisement messages in this case. Click the **Settings** tab and click the **IPv6** tab and then uncheck the **Enable RA** check box to suppress messages.

- Manual configuration—To manually configure a global IPv6 address:

  a. Click the **Address** page, and click (+)**Add Address**.

The **Add Address** dialog box appears.

b. In the **Address** field, enter either a full global IPv6 address, including the interface ID, or enter the IPv6 prefix, along with the IPv6 prefix length. (Routed Mode) If you only enter the prefix, then be sure to check the **Enforce EUI 64** check box to generate the interface ID using the Modified EUI-64 format. For example, 2001:0DB8::BA98:0:3210/48 (full address) or 2001:0DB8::/48 (prefix, with EUI 64 checked).

For High Availability (if you did not set **Enforce EUI 64**), set the standby IP address on the **Devices** > **Device Management** page, click **High Availability** in the **Monitored Interfaces** area. If you do not set the standby IP address, the active unit cannot monitor the standby interface using network tests; it can only track the link state.

- (Routed interface) Obtain an address using DHCPv6—To use DHCPv6:

Figure 379: Enable the DHCPv6 Client

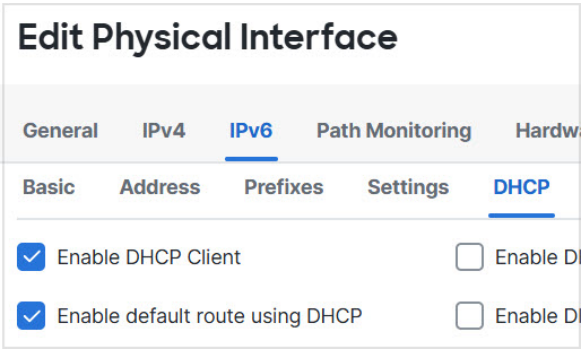

a. Click the **DHCP** page.

b. Check the check box of **Enable DHCP Client**.

c. (Optional) Check the check box of **Enable default route using DHCP** to obtain a default route from Router Advertisements.

- (Routed interface) Use a delegated prefix—To assign an IPv6 address using the delegated prefix:

This feature requires the Firewall Threat Defense to have the DHCPv6 Prefix Delegation client enabled *on a different interface*. See Enable the IPv6 Prefix Delegation Client, on page 867.

a. Click the **DHCP** page.

b. Click (+).

Figure 380: Use a Delegated Prefix

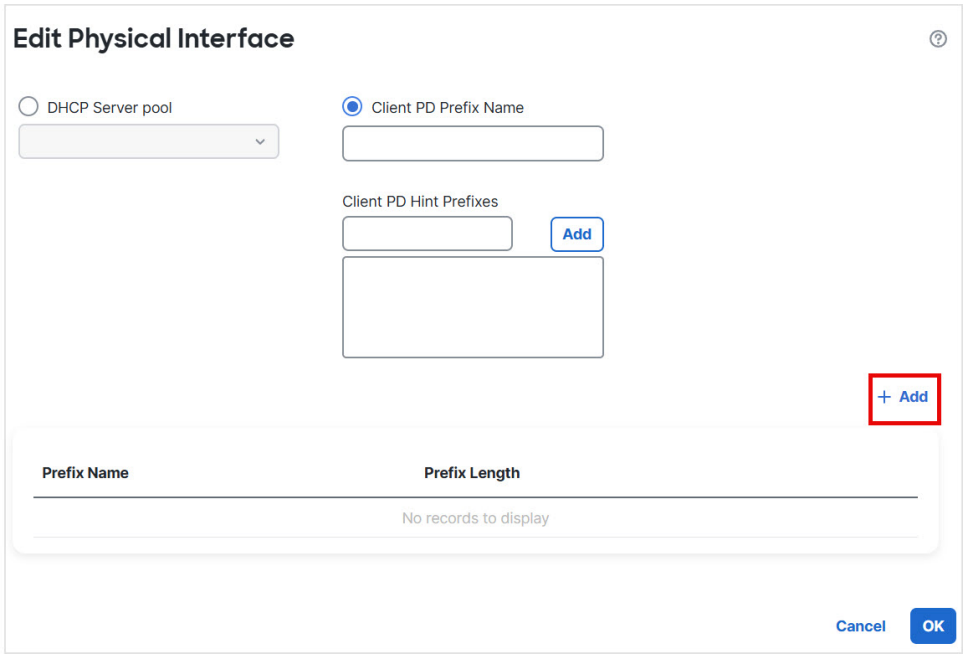

c. Enter the **Prefix Name** that you specified for the Prefix Delegation client (see Enable the IPv6 Prefix Delegation Client, on page 867) on another interface.

Figure 381: Specify the Prefix Name and Address

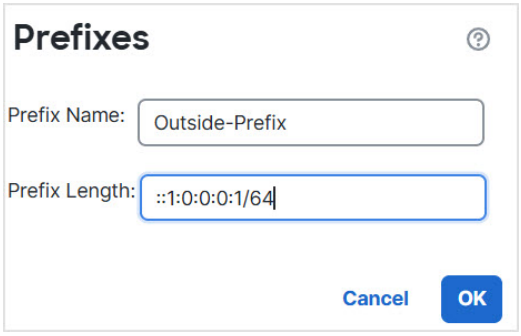

d. Enter the IPv6 address and **Prefix Length**.

Typically, the delegated prefix will be /60 or smaller so you can subnet to multiple /64 networks. /64 is the supported subnet length if you want to support SLAAC for connected clients. You should specify an address that completes the /60 subnet, for example ::1:0:0:0:1. Enter :: before the address in case the prefix is smaller than /60. For example, if the delegated prefix is 2001:DB8:1234:5670::/60, then the global IP address assigned to this interface is 2001:DB8:1234:5671::1/64. The prefix that is advertised in router advertisements is 2001:DB8:1234:5671::/64. In this example, if the prefix is smaller than /60, the remaining bits of the prefix will be 0's as indicated by the leading ::. For example, if the prefix is 2001:DB8:1234::/48, then the IPv6 address will be 2001:DB8:1234::1:0:0:0:1/64.

e. Click **OK**.

Figure 382: Prefix Delegation Table

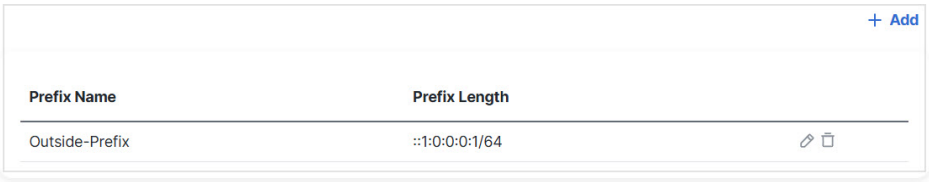

f. Optionally enable the DHCPv6 stateless server on this interface (see Enable the DHCPv6 Stateless Server, on page 917). If you do so, we recommend that you also check the **Enable DHCP for non-address config** option.

Step 6 For Routed interfaces, you can optionally set the following values on the **Basic** page:

- To enforce the use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link, check the **Enforce EUI-64** check box.
- To manually set the link-local address, enter an address in the **Link-Local address** field.

  A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. If you do not want to configure a global address, and only need to configure a link-local address, you have the option of manually defining the link-local address. Note that we recommend automatically assigning the link-local address based on the Modified EUI-64 format. For example, if other devices enforce the use of the Modified EUI-64 format, then a manually-assigned link-local address may cause packets to be dropped.

  Clustering does not support manual link-local addresses.

Step 7 For Routed interfaces, you can optionally set the following values on the **DHCP** page:

- Check the **Enable DHCP for address config** check box to set the Managed Address Config flag in the IPv6 router advertisement packet.

  This flag in IPv6 router advertisements informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain addresses, in addition to the derived stateless autoconfiguration address.
- Check the **Enable DHCP for non-address config** check box to set the Other Address Config flag in the IPv6 router advertisement packet.

  This flag in IPv6 router advertisements informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain additional information from DHCPv6, such as the DNS server address. Use this option when using the DHCPv6 stateless server with DHCPv6 prefix delegation.

Step 8 For Routed interfaces, see Configure IPv6 Neighbor Discovery, on page 873 to configure settings on the **Prefixes** and **Settings** pages. For BVI interfaces, see the following parameters on the **Settings** page:

- **DAD attempts**—The maximum number of DAD attempts, between 1 and 600. Set the value to 0 to disable duplicate address detection (DAD) processing. This setting configures the number of consecutive neighbor solicitation messages that are sent on an interface while DAD is performed on IPv6 addresses. 1 attempt is the default.
- **NS Interval**—The interval between IPv6 neighbor solicitation retransmissions on an interface, between 1000 and 3600000 ms. The default value is 1000 ms.

- **Reachable Time**—The amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event has occurred, between 0 and 3600000 ms. The default value is 0 ms. When 0 is used for the value, the reachable time is sent as undetermined. It is up to the receiving devices to set and track the reachable time value. The neighbor reachable time enables detecting unavailable neighbors. Shorter configured times enable detecting unavailable neighbors more quickly, however, shorter times consume more IPv6 network bandwidth and processing resources in all IPv6 network devices. Very short configured times are not recommended in normal IPv6 operation.

Step 9 Click **OK**.

Step 10 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure IPv6 Neighbor Discovery

The IPv6 neighbor discovery process uses ICMPv6 messages and solicited-node multicast addresses to determine the link-layer address of a neighbor on the same network (local link), verify the readability of a neighbor, and keep track of neighboring routers.

Nodes (hosts) use neighbor discovery to determine the link-layer addresses for neighbors known to reside on attached links and to quickly purge cached values that become invalid. Hosts also use neighbor discovery to find neighboring routers that are willing to forward packets on their behalf. In addition, nodes use the protocol to actively keep track of which neighbors are reachable and which are not, and to detect changed link-layer addresses. When a router or the path to a router fails, a host actively searches for functioning alternates.

Before you begin

Supported in Routed mode only. For IPv6 neighbor settings supported in transparent mode, see Configure a Global IPv6 Address, on page 868.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click **IPv6**, and then **Prefixes**.

Step 4 (Optional) To configure which IPv6 prefixes are included in IPv6 router advertisements, perform the following steps:

a) Click (+)**Add Prefix**.

b) In the **Address** field, enter the IPv6 address with the prefix length or check the **Default** check box to use the default prefix.

c) (Optional) Uncheck the **Advertisement** check box to indicate that the IPv6 prefix is not advertised. For the **Default** prefix, this setting only applies to on-link prefixes. Off-link prefixes will still be advertised unless you uncheck **Advertisement** for a specific off-link prefix.

d) Check the **Off Link** check box to indicate that the specified prefix is assigned to the link. Nodes sending traffic to addresses that contain the specified prefix consider the destination to be locally reachable on the link. This prefix should not be used for on-link determination.

e) To use the specified prefix for autoconfiguration, check the **Autoconfiguration** check box.

f) For the **Prefix Lifetime**, click **Duration** or **Expiration Date**.

- **Duration**—Enter a **Preferred Lifetime** for the prefix in seconds. This setting is the amount of time that the specified IPv6 prefix is advertised as being valid. The maximum value represents infinity. Valid values are from 0 to 4294967295. The default is 2592000 (30 days). Enter a **Valid Lifetime** for the prefix in seconds. This setting is the amount of time that the specified IPv6 prefix is advertised as being preferred. The maximum value represents infinity. Valid values are from 0 to 4294967295. The default setting is 604800 (seven days). Alternatively, check the **Infinite** check box to set an unlimited duration.
- **Expiration Date**—Choose a **Valid** and **Preferred** date and time.

g) Click **OK**.

Step 5 Click **Settings**.

Step 6 (Optional) Set the maximum number of **DAD attempts**, between 1 and 600. 1 attempt is the default. Set the value to 0 to disable duplicate address detection (DAD) processing.

This setting configures the number of consecutive neighbor solicitation messages that are sent on an interface while DAD is performed on IPv6 addresses.

During the stateless autoconfiguration process, Duplicate Address Detection verifies the uniqueness of new unicast IPv6 addresses before the addresses are assigned to interfaces.

When a duplicate address is identified, the state of the address is set to DUPLICATE, the address is not used, and the following error message is generated:

```
325002: Duplicate address ipv6_address/MAC_address on interface
```

If the duplicate address is the link-local address of the interface, the processing of IPv6 packets is disabled on the interface. If the duplicate address is a global address, the address is not used.

Step 7 (Optional) Configure the interval between IPv6 neighbor solicitation retransmissions in the **NS Interval** field, between 1000 and 3600000 ms.

The default value is 1000 ms.

Neighbor solicitation messages (ICMPv6 Type 135) are sent on the local link by nodes attempting to discover the link-layer addresses of other nodes on the local link. After receiving a neighbor solicitation message, the destination node replies by sending a neighbor advertisement message (ICPMv6 Type 136) on the local link.

After the source node receives the neighbor advertisement, the source node and destination node can communicate. Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. When a node wants to verifying the reachability of a neighbor, the destination address in a neighbor solicitation message is the unicast address of the neighbor.

Neighbor advertisement messages are also sent when there is a change in the link-layer address of a node on a local link.

Step 8 (Optional) Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event has occurred in the **Reachable Time** field, between 0 and 3600000 ms.

The default value is 0 ms. When 0 is used for the value, the reachable time is sent as undetermined. It is up to the receiving devices to set and track the reachable time value.

The neighbor reachable time enables detecting unavailable neighbors. Shorter configured times enable detecting unavailable neighbors more quickly, however, shorter times consume more IPv6 network bandwidth and processing resources in all IPv6 network devices. Very short configured times are not recommended in normal IPv6 operation.

Step 9 (Optional) To suppress the router advertisement transmissions, uncheck the **Enable RA** check box. If you enable router advertisement transmissions, you can set the RA lifetime and interval.

Router advertisement messages (ICMPv6 Type 134) are automatically sent in response to router solicitation messages (ICMPv6 Type 133). Router solicitation messages are sent by hosts at system startup so that the host can immediately autoconfigure without needing to wait for the next scheduled router advertisement message.

You may want to disable these messages on any interface for which you do not want the Firewall Threat Defense to supply the IPv6 prefix (for example, the outside interface).

- **RA Lifetime**—Configure the router lifetime value in IPv6 router advertisements, between 0 and 9000 seconds.

  The default is 1800 seconds.

- **RA Interval**—Configure the interval between IPv6 router advertisement transmissions, between 3 and 1800 seconds.

  The default is 200 seconds.

  To prevent synchronization with other IPv6 nodes, the firewall randomly adjusts the value that you set (jitter).

Step 10 Click **OK**.

Step 11 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Advanced Interface Settings

This section describes how to configure MAC addresses for regular firewall mode interfaces, how to set the maximum transmission unit (MTU), and how to set other advanced parameters.

About Advanced Interface Configuration

This section describes advanced interface settings.

About MAC Addresses

You can manually assign MAC addresses to override the default. For container instances, the FXOS chassis automatically generates unique MAC addresses for all interfaces.

Note You might want to assign unique MAC addresses to subinterfaces defined on the Firewall Threat Defense, because they use the same burned-in MAC address of the parent interface. For example, your service provider might perform access control based on the MAC address. Also, because IPv6 link-local addresses are generated based on the MAC address, assigning unique MAC addresses to subinterfaces allows for unique IPv6 link-local addresses, which can avoid traffic disruption in certain instances on the Firewall Threat Defense device.

Note For container instances, even if you are not sharing a subinterface, if you manually configure MAC addresses, make sure you use unique MAC addresses for all subinterfaces on the same parent interface to ensure proper classification.

Default MAC Addresses

For native instances:

Default MAC address assignments depend on the type of interface.

- Physical interfaces—The physical interface uses the burned-in MAC address.
- VLAN interfaces (Firepower 1010 and Secure Firewall 1210/1220)—Routed firewall mode: All VLAN interfaces share a MAC address. Ensure that any connected switches can support this scenario. If the connected switches require unique MAC addresses, you can manually assign MAC addresses. See Configure the MAC Address, on page 881.

  Transparent firewall mode: Each VLAN interface has a unique MAC address. You can override the generated MAC addresses if desired by manually assigning MAC addresses. See Configure the MAC Address, on page 881.
- EtherChannels (Firepower Models)—For an EtherChannel, all interfaces that are part of the channel group share the same MAC address. This feature makes the EtherChannel transparent to network applications and users, because they only see the one logical connection; they have no knowledge of the individual links. The port-channel interface uses a unique MAC address from a pool; interface membership does not affect the MAC address.
- EtherChannels (ASA Models)—The port-channel interface uses the lowest-numbered channel group interface MAC address as the port-channel MAC address. Alternatively you can configure a MAC address for the port-channel interface. We recommend configuring a unique MAC address in case the group channel interface membership changes. If you remove the interface that was providing the port-channel MAC address, then the port-channel MAC address changes to the next lowest numbered interface, thus causing traffic disruption.
- Subinterfaces (Firewall Threat Defense-defined)—All subinterfaces of a physical interface use the same burned-in MAC address. You might want to assign unique MAC addresses to subinterfaces. For example, your service provider might perform access control based on the MAC address. Also, because IPv6 link-local addresses are generated based on the MAC address, assigning unique MAC addresses to subinterfaces allows for unique IPv6 link-local addresses, which can avoid traffic disruption in certain instances on the Firewall Threat Defense.

For container instances:

- MAC addresses for all interfaces are taken from a MAC address pool. For subinterfaces, if you decide to manually configure MAC addresses, make sure you use unique MAC addresses for all subinterfaces on the same parent interface to ensure proper classification. See Automatic MAC Addresses for Container Instance Interfaces, on page 319.

About the MTU

The MTU specifies the maximum frame *payload* size that the Firewall Threat Defense device can transmit on a given Ethernet interface. The MTU value is the frame size *without* Ethernet headers, VLAN tagging, or other overhead. For example, when you set the MTU to 1500, the expected frame size is 1518 bytes including the headers, or 1522 when using VLAN. Do not set the MTU value higher to accommodate these headers.

For Geneve, the entire Ethernet datagram is being encapsulated, so the new IP packet is larger and requires a larger MTU: you should set the Firewall Threat Defense device VTEP source interface MTU to be the network MTU + 306 bytes.

Path MTU Discovery

The Firewall Threat Defense device supports Path MTU Discovery (as defined in RFC 1191), which lets all devices in a network path between two hosts coordinate the MTU so they can standardize on the lowest MTU in the path.

Default MTU

The default MTU on the Firewall Threat Defense device is 1500 bytes. This value does not include the 18-22 bytes for the Ethernet header, VLAN tagging, or other overhead.

MTU and Fragmentation

For IPv4, if an outgoing IP packet is larger than the specified MTU, it is fragmented into 2 or more frames. Fragments are reassembled at the destination (and sometimes at intermediate hops), and fragmentation can cause performance degradation. For IPv6, packets are typically not allowed to be fragmented at all. Therefore, your IP packets should fit within the MTU size to avoid fragmentation.

For TCP packets, the endpoints typically use their MTU to determine the TCP maximum segment size (MTU - 40, for example). If additional TCP headers are added along the way, for example for site-to-site VPN tunnels, then the TCP MSS might need to be adjusted down by the tunneling entity. See About the TCP MSS, on page 878.

For UDP or ICMP, the application should take the MTU into account to avoid fragmentation.

Note The Firewall Threat Defense device can receive frames larger than the configured MTU as long as there is room in memory.

MTU and Jumbo Frames

A larger MTU lets you send larger packets. Larger packets might be more efficient for your network. See the following guidelines:

- Matching MTUs on the traffic path—We recommend that you set the MTU on all Firewall Threat Defense interfaces and other device interfaces along the traffic path to be the same. Matching MTUs prevents intermediate devices from fragmenting the packets.

- Accommodating jumbo frames—You can set the MTU 9000 bytes or higher when you enable jumbo frames. The maximum depends on the model.

About the TCP MSS

The TCP maximum segment size (MSS) is the size of the TCP payload *before* any TCP and IP headers are added. UDP packets are not affected. The client and the server exchange TCP MSS values during the three-way handshake when establishing the connection.

You can set the TCP MSS on the Firewall Threat Defense device for through traffic using the Sysopt_Basic object in FlexConfig; see FlexConfig Policies, on page 2667; by default, the maximum TCP MSS is set to 1380 bytes. This setting is useful when the Firewall Threat Defense device needs to add to the size of the packet for IPsec VPN encapsulation. However, for non-IPsec endpoints, you should disable the maximum TCP MSS on the Firewall Threat Defense device.

If you set a maximum TCP MSS, if either endpoint of a connection requests a TCP MSS that is larger than the value set on the Firewall Threat Defense device, then the Firewall Threat Defense device overwrites the TCP MSS in the request packet with the Firewall Threat Defense device maximum. If the host or server does not request a TCP MSS, then the Firewall Threat Defense device assumes the RFC 793-default value of 536 bytes (IPv4) or 1220 bytes (IPv6), but does not modify the packet. For example, you leave the default MTU as 1500 bytes. A host requests an MSS of 1500 minus the TCP and IP header length, which sets the MSS to 1460. If the Firewall Threat Defense device maximum TCP MSS is 1380 (the default), then the Firewall Threat Defense device changes the MSS value in the TCP request packet to 1380. The server then sends packets with 1380-byte payloads. The Firewall Threat Defense device can then add up to 120 bytes of headers to the packet and still fit in the MTU size of 1500.

You can also configure the minimum TCP MSS; if a host or server requests a very small TCP MSS, the Firewall Threat Defense device can adjust the value up. By default, the minimum TCP MSS is not enabled.

For to-the-box traffic, including for SSL VPN connections, this setting does not apply. The Firewall Threat Defense device uses the MTU to derive the TCP MSS: MTU - 40 (IPv4) or MTU - 60 (IPv6).

Default TCP MSS

By default, the maximum TCP MSS on the Firewall Threat Defense device is 1380 bytes. This default accommodates IPv4 IPsec VPN connections where the headers can equal up to 120 bytes; this value fits within the default MTU of 1500 bytes.

Suggested Maximum TCP MSS Setting

The default TCP MSS assumes the Firewall Threat Defense device acts as an IPv4 IPsec VPN endpoint and has an MTU of 1500. When the Firewall Threat Defense device acts as an IPv4 IPsec VPN endpoint, it needs to accommodate up to 120 bytes for TCP and IP headers.

If you change the MTU value, use IPv6, or do not use the Firewall Threat Defense device as an IPsec VPN endpoint, then you should change the TCP MSS setting using the Sysopt_Basic object in FlexConfig.

Note Even if you explicitly set an MSS, if a component such as TLS/SSL decryption or server discovery needs a particular MSS, it will set that MSS based on the interface MTU and ignore your MSS setting.

See the following guidelines:

- Normal traffic—Disable the TCP MSS limit and accept the value established between connection endpoints. Because connection endpoints typically derive the TCP MSS from the MTU, non-IPsec packets usually fit this TCP MSS.
- IPv4 IPsec endpoint traffic—Set the maximum TCP MSS to the MTU - 120. For example, if you use jumbo frames and set the MTU to 9000, then you need to set the TCP MSS to 8880 to take advantage of the new MTU.
- IPv6 IPsec endpoint traffic—Set the maximum TCP MSS to the MTU - 140.

ARP Inspection for Bridge Group Traffic

By default, all ARP packets are allowed between bridge group members. You can control the flow of ARP packets by enabling ARP inspection.

ARP inspection prevents malicious users from impersonating other hosts or routers (known as ARP spoofing). ARP spoofing can enable a "man-in-the-middle" attack. For example, a host sends an ARP request to the gateway router; the gateway router responds with the gateway router MAC address. The attacker, however, sends another ARP response to the host with the attacker MAC address instead of the router MAC address. The attacker can now intercept all the host traffic before forwarding it on to the router.

ARP inspection ensures that an attacker cannot send an ARP response with the attacker MAC address, so long as the correct MAC address and the associated IP address are in the static ARP table.

When you enable ARP inspection, the Firewall Threat Defense device compares the MAC address, IP address, and source interface in all ARP packets to static entries in the ARP table, and takes the following actions:

- If the IP address, MAC address, and source interface match an ARP entry, the packet is passed through.
- If there is a mismatch between the MAC address, the IP address, or the interface, then the Firewall Threat Defense device drops the packet.
- If the ARP packet does not match any entries in the static ARP table, then you can set the Firewall Threat Defense device to either forward the packet out all interfaces (flood), or to drop the packet.

Note The dedicated Management interface never floods packets even if this parameter is set to flood.

MAC Address Table

When you use bridge groups, the Firewall Threat Defense learns and builds a MAC address table in a similar way as a normal bridge or switch: when a device sends a packet through the bridge group, the Firewall Threat Defense adds the MAC address to its table. The table associates the MAC address with the source interface so that the Firewall Threat Defense knows to send any packets addressed to the device out the correct interface. Because traffic between bridge group members is subject to the Firewall Threat Defense security policy, if the destination MAC address of a packet is not in the table, the Firewall Threat Defense does not flood the original packet on all interfaces as a normal bridge does. Instead, it generates the following packets for directly-connected devices or for remote devices:

- Packets for directly-connected devices—The Firewall Threat Defense generates an ARP request for the destination IP address, so that it can learn which interface receives the ARP response.

- Packets for remote devices—The Firewall Threat Defense generates a ping to the destination IP address so that it can learn which interface receives the ping reply.

The original packet is dropped.

Default Settings

- If you enable ARP inspection, the default setting is to flood non-matching packets.
- The default timeout value for dynamic MAC address table entries is 5 minutes.
- By default, each interface automatically learns the MAC addresses of entering traffic, and the Firewall Threat Defense device adds corresponding entries to the MAC address table.

Note Secure Firewall Threat Defense device generates a reset packet to reset a connection that is denied by a stateful inspection engine. Here, the destination MAC address of the packet is not determined based on the ARP table lookup but instead it is taken directly from the packets (connections) that are being denied.

Guidelines for ARP Inspection and the MAC Address Table

- ARP inspection is only supported for bridge groups.
- MAC address table configuration is only supported for bridge groups.

Configure the MTU

Customize the MTU on the interface, for example, to allow jumbo frames.

For the ISA 3000 and the Firewall Threat Defense Virtual: Changing the MTU above 1500 bytes automatically enables jumbo-frame reservation. You must restart the system before you can use jumbo frames. For the Firewall Threat Defense Virtual that supports clustering, you can enable jumbo-frame reservation in the Day0 configuration, so in that case, you do not need to restart. After you restart, you cannot disable jumbo-frame reservation. An exception is for the Firewall Threat Defense Virtual, where you can disable jumbo-frame reservation in the Day0 configuration, if supported. If you use an interface in an inline set, the MTU setting is not used. However, the jumbo-frame reservation setting *is* relevant to inline sets; jumbo frames enable the inline interfaces to receive packets up to 9000 bytes. To enable jumbo-frame reservation, you must set the MTU of *any* interface above 1500 bytes.

Jumbo frames are enabled by default on other platforms.

Caution Changing the highest MTU value on the device for a data interface restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all data interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. This caution does not apply to management-only interfaces. See Snort Restart Traffic Behavior, on page 254 for more information.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 On the **General** tab, set the **MTU**. The minimum and maximum depends on your platform.

The default is 1500 bytes.

Step 4 Click **OK**.

Step 5 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Step 6 For the ISA 3000 and the Firewall Threat Defense Virtual, if you set the MTU above 1500 bytes, restart the system to enable jumbo-frame reservation. See Shut Down or Restart the Device, on page 62.

Configure the MAC Address

You might need to manually assign a MAC address. You can also set the Active and Standby MAC addresses on the **Devices** > **Device Management** window by choosing **High Availability** from the **Add** drop-down list. If you set the MAC address for an interface on both screens, the addresses on the **Interfaces** > **Advanced** tab take precedence.

Note For container instances, even if you are not sharing a subinterface, if you manually configure MAC addresses, make sure you use unique MAC addresses for all subinterfaces on the same parent interface to ensure proper classification.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click the **Advanced** tab.
The **Information** tab is selected.

Step 4 (For clustering in individual interface mode) Choose a MAC address pool from the drop-down list.

You can add MAC address pools according to Address Pools, on page 1374.

Step 5 (For other modes) Set the active and standby MAC addresses.

a) In the **Active MAC Address** field, enter a MAC address in H.H.H format, where H is a 16-bit hexadecimal digit.

For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE. The MAC address must not have the multicast bit set, that is, the second hexadecimal digit from the left cannot be an odd number.

b) In the **Standby MAC Address** field, enter a MAC address for use with High Availability.

If the active unit fails over and the standby unit becomes active, the new active unit starts using the active MAC addresses to minimize network disruption, while the old active unit uses the standby address.

Step 6 Click **OK**.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Add a Static ARP Entry

By default, all ARP packets are allowed between bridge group members. You can control the flow of ARP packets by enabling ARP inspection (see ARP Inspection, on page 947). ARP inspection compares ARP packets with *static* ARP entries in the ARP table.

For routed interfaces, you can enter static ARP entries, but normally dynamic entries are sufficient. For routed interfaces, the ARP table is used to deliver packets to directly-connected hosts. Although senders identify a packet destination by an IP address, the actual delivery of the packet on Ethernet relies on the Ethernet MAC address. When a router or host wants to deliver a packet on a directly connected network, it sends an ARP request asking for the MAC address associated with the IP address, and then delivers the packet to the MAC address according to the ARP response. The host or router keeps an ARP table so it does not have to send ARP requests for every packet it needs to deliver. The ARP table is dynamically updated whenever ARP responses are sent on the network, and if an entry is not used for a period of time, it times out. If an entry is incorrect (for example, the MAC address changes for a given IP address), the entry needs to time out before it can be updated with the new information.

For transparent mode, the Firewall Threat Defense only uses dynamic ARP entries in the ARP table for traffic to and from the Firewall Threat Defense device, such as management traffic.

Before you begin

This screen is only available for named interfaces.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click the **Advanced** tab, and then click the **ARP** tab (called **ARP and MAC** for transparent mode).

Step 4 Click (+)**Add ARP Config**.
The **Add ARP Config** dialog box appears.

Step 5 In the **IP Address** field, enter the IP address of the host.

Step 6 In the **MAC Address** field, enter the MAC address of the host; for example, 00e0.1e4e.3d8b.

Step 7 To perform proxy ARP for this address, check the **Enable Alias** check box.

If the Firewall Threat Defense device receives an ARP request for the specified IP address, then it responds with the specified MAC address.

Step 8 Click **OK**, and then click **OK** again to exit the Advanced settings.

Step 9 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Add a Static MAC Address and Disable MAC Learning for a Bridge Group

Normally, MAC addresses are added to the MAC address table dynamically as traffic from a particular MAC address enters an interface. You can disable MAC address learning; however, unless you statically add MAC addresses to the table, no traffic can pass through the Firewall Threat Defense device. You can also add static MAC addresses to the MAC address table. One benefit to adding static entries is to guard against MAC spoofing. If a client with the same MAC address as a static entry attempts to send traffic to an interface that does not match the static entry, then the Firewall Threat Defense device drops the traffic and generates a system message. When you add a static ARP entry (see Add a Static ARP Entry, on page 882), a static MAC address entry is automatically added to the MAC address table.

Before you begin

This screen is only available for named BVIs in transparent mode.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click the **Advanced** tab, and then click the **ARP and MAC** tab.

Step 4 (Optional) Disable MAC learning by unchecking the **Enable MAC Learning** check box.

Step 5 To add a static MAC address, click **Add MAC Config**.
The **Add MAC Config** dialog box appears.

Step 6 In the **MAC Address** field, enter the MAC address of the host; for example, 00e0.1e4e.3d8b. Click **OK**.

Step 7 Click **OK** to exit the Advanced settings.

Step 8 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Set Security Configuration Parameters

This section describes how to prevent IP spoofing, allow full fragment reassembly, and override the default fragment setting set for at the device level in **Platform Settings** .

Anti-Spoofing

This section lets you enable Unicast Reverse Path Forwarding on an interface. Unicast RPF guards against IP spoofing (a packet uses an incorrect source IP address to obscure its true source) by ensuring that all packets have a source IP address that matches the correct source interface according to the routing table.

Normally, the Firewall Threat Defense device only looks at the destination address when determining where to forward the packet. Unicast RPF instructs the device to also look at the source address; this is why it is called Reverse Path Forwarding. For any traffic that you want to allow through the Firewall Threat Defense device, the device routing table must include a route back to the source address. See RFC 2267 for more information.

For outside traffic, for example, the Firewall Threat Defense device can use the default route to satisfy the Unicast RPF protection. If traffic enters from an outside interface, and the source address is not known to the routing table, the device uses the default route to correctly identify the outside interface as the source interface.

If traffic enters the outside interface from an address that is known to the routing table, but is associated with the inside interface, then the Firewall Threat Defense device drops the packet. Similarly, if traffic enters the inside interface from an unknown source address, the device drops the packet because the matching route (the default route) indicates the outside interface.

Unicast RPF is implemented as follows:

- ICMP packets have no session, so each packet is checked.
- UDP and TCP have sessions, so the initial packet requires a reverse route lookup. Subsequent packets arriving during the session are checked using an existing state maintained as part of the session. Non-initial packets are checked to ensure they arrived on the same interface used by the initial packet.

Fragment per Packet

By default, the Firewall Threat Defense device allows up to 24 fragments per IP packet, and up to 200 fragments awaiting reassembly. You might need to let fragments on your network if you have an application that routinely fragments packets, such as NFS over UDP. However, if you do not have an application that fragments traffic, we recommend that you do not allow fragments through the Firewall Threat Defense device. Fragmented packets are often used as DoS attacks.

Fragment Reassembly

The Firewall Threat Defense device performs the following fragment reassembly processes:

- IP fragments are collected until a fragment set is formed or until a timeout interval has elapsed.
- If a fragment set is formed, integrity checks are performed on the set. These checks include no overlapping, no tail overflow, and no chain overflow.
- IP fragments that terminate at the Firewall Threat Defense device are always fully reassembled.
- If **Full Fragment Reassembly** is disabled (the default), the fragment set is forwarded to the transport layer for further processing.
- If **Full Fragment Reassembly** is enabled, the fragment set is first coalesced into a single IP packet. The single IP packet is then forwarded to the transport layer for further processing.

Before you begin

This screen is only available for named interfaces.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click the **Advanced** tab, and then click the **Security Configuration** tab.

Step 4 To enable Unicast Reverse Path Forwarding, check the **Enable Anti Spoofing** check box.

Step 5 To enable full fragment reassembly, check the **Allow Full Fragment Reassembly** check box.

Step 6 To change the number of fragments allowed per packet, check the **Override Default Fragment Setting** check box, and set the following values:

- **Size**—Set the maximum number of packets that can be in the IP reassembly database waiting for reassembly. The default is 200. Set this value to 1 to disable fragments.
- **Chain**—Set the maximum number of packets into which a full IP packet can be fragmented. The default is 24 packets.
- **Timeout**—Set the maximum number of seconds to wait for an entire fragmented packet to arrive. The timer starts after the first fragment of a packet arrives. If all fragments of the packet do not arrive by the number of seconds specified, all fragments of the packet that were already received will be discarded. The default is 5 seconds.

Step 7 Click **OK**.

Step 8 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

History for Regular Firewall Interfaces

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Secure Firewall 1210CP IEEE 802.3bt support (PoE++ and Hi-PoE)	7.7.0	7.7.0	See the following improvements related to support for IEEE 802.3bt: • PoE++ and Hi-PoE—Up to 90W per port. • Single- and dual-signature powered devices (PDs). • Power budgeting is done on a first-come, first-served basis. • Power budget fields were added to **show power inline**. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** > **PoE** New/Modified commands: **show power inline**
Secure Firewall supports Dual-Arm Deployment Mode on AWS with GWLB	7.6	7.6	The Secure Firewall supports the dual-arm deployment mode on AWS with GWLB. This mode enables the firewall to directly forward internet-bound traffic to the internet through the internet gateway after traffic inspection, while also performing network address translation (NAT).
Secure Firewall 1210/1220 hardware switch support	7.6	7.6	The Secure Firewall 1210/1220 supports setting each Ethernet interface to be a switch port or a firewall interface
Secure Firewall 1210CP PoE+ support on Ethernet ports 1/5-1/8	7.6	7.6	The Secure Firewall 1210CP supports Power over Ethernet+ (PoE+) on Ethernet ports 1/5-1/8.
VXLAN VTEP IPv6 support	7.4	Any	You can now specify an IPv6 address for the VXLAN VTEP interface. IPv6 is not supported for the threat defense virtual cluster control link or for Geneve encapsulation. New/Modified screens: • **Devices** > **Device Management** > **Edit** > **VTEP** > **Add VTEP** **Devices** > **Device Management** > **Edit** > **Interfaces** > **Add Interfaces** > **VNI Interface** Requires Firewall Threat Defense version 7.4.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Loopback interface support for BGP and management traffic	7.4	Any	You can use a loopback interface for: • AAA • BGP • DNS • HTTP • ICMP • IPsec Flow Offload • NetFlow • SNMP • SSH • Syslog Requires Firewall Threat Defense version 7.4.
Loopback interface support for VTI	7.3	Any	You can now add a loopback interface. The loopback interface helps to overcome path failures. If an interface goes down, you can access all interfaces through the IP address assigned to the loopback interface. For VTI, in addition to setting a loopback interface as the source interface, support has also been added to inherit the IP address from a loopback interface instead of a statically configured IP address. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Add Interfaces** > **Add Loopback Interface**

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
IPv6 DHCP	7.3	Any	The Firewall Threat Defense now supports the following features for IPv6 addressing: • DHCPv6 Address client—The Firewall Threat Defense obtains an IPv6 global address and optional default route from the DHCPv6 server. • DHCPv6 Prefix Delegation client—The Firewall Threat Defense obtains delegated prefix(es) from a DHCPv6 server. The Firewall Threat Defense can then use these prefixes to configure other Firewall Threat Defense interface addresses so that StateLess Address Auto Configuration (SLAAC) clients can autoconfigure IPv6 addresses on the same network. • BGP router advertisement for delegated prefixes • DHCPv6 stateless server—The Firewall Threat Defense provides other information such as the domain name to SLAAC clients when they send Information Request (IR) packets to the Firewall Threat Defense. The Firewall Threat Defense only accepts IR packets, and does not assign addresses to the clients. New/Modified screens: • **Devices** > **Device Management** > **Interfaces** > **Add/Edit Interfaces** > **IPv6** > **DHCP** • **Objects** > **Object Management** > **DHCP IPv6 Pool** New/Modified commands: **show bgp ipv6 unicast, show ipv6 dhcp, show ipv6 general-prefix**
Paired proxy VXLAN for the Firewall Threat Defense Virtual for the Azure Gateway Load Balancer	7.3	Any	You can configure a paired proxy mode VXLAN interface for the Firewall Threat Defense Virtual in Azure for use with the Azure Gateway Load Balancer (GWLB). The Firewall Threat Defense Virtual defines an external interface and an internal interface on a single NIC by utilizing VXLAN segments in a paired proxy. New/Modified screens: • **Devices** > **Device Management** > **Device** > **Interfaces** > **Add Interfaces** > **VNI Interface** Supported platforms: Firewall Threat Defense Virtual in Azure

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
VXLAN support	7.2	Any	VXLAN encapsulation support was added. New/Modified screens: • **Devices** > **Device Management** > **Device** > **VTEP** • **Devices** > **Device Management** > **Device** > **Interfaces** > **Add Interfaces** > **VNI Interface** • **Devices** > **Device Management** > **Device** > **Interfaces** edit physical interface **> General** Supported platforms: All.
Geneve support for the Firewall Threat Defense Virtual	7.1	Any	Geneve encapsulation support was added for the Firewall Threat Defense Virtual to support single-arm proxy for the Amazon Web Services (AWS) Gateway Load Balancer. The AWS Gateway Load Balancer combines a transparent network gateway (with a single entry and exit point for all traffic) and a load balancer that distributes traffic and scales Firewall Threat Defense Virtual to match the traffic demand. This feature requires Snort 3. New/Modified screens: • **Devices** > **Device Management** > **Device** > **VTEP** • **Devices** > **Device Management** > **Device** > **Interfaces** > **Add Interfaces** > **VNI Interface** • **Devices** > **Device Management** > **Device** > **Interfaces** edit physical interface **> General** Supported platforms: Firewall Threat Defense Virtual in AWS
31-bit Subnet Mask	7.0	Any	For routed interfaces, you can configure an IP address on a 31-bit subnet for point-to-point connections. The 31-bit subnet includes only 2 addresses; normally, the first and last address in the subnet is reserved for the network and broadcast, so a 2-address subnet is not usable. However, if you have a point-to-point connection and do not need network or broadcast addresses, a 31-bit subnet is a useful way to preserve addresses in IPv4. For example, the failover link between 2 FTDs only requires 2 addresses; any packet that is transmitted by one end of the link is always received by the other, and broadcasting is unnecessary. You can also have a directly-connected management station running SNMP or Syslog. This feature is not supported for BVIs for bridge groups or with multicast routing. New/Modified screens: **Devices** > **Device Management** > **Interfaces**

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Synchronization between the Firewall Threat Defense operational link state and the physical link state for the Firepower 4100/9300	6.7	Any	The Firepower 4100/9300 chassis can now synchronize the Firewall Threat Defense operational link state with the physical link state for data interfaces. Currently, interfaces will be in an Up state as long as the FXOS admin state is up and the physical link state is up. The Firewall Threat Defense application interface admin state is not considered. Without synchronization from Firewall Threat Defense, data interfaces can be in an Up state physically before the Firewall Threat Defense application has completely come online, for example, or can stay Up for a period of time after you initiate an Firewall Threat Defense shutdown. For inline sets, this state mismatch can result in dropped packets because external routers may start sending traffic to the Firewall Threat Defense before the Firewall Threat Defense can handle it. This feature is disabled by default, and can be enabled per logical device in FXOS. **Note** This feature is not supported for clustering, container instances, or Firewall Threat Defense with a Radware vDP decorator. It is also not supported for ASA. New/Modified Firepower Chassis Manager screens: **Logical Devices** > **Enable Link State** New/Modified FXOS commands: **set link-state-sync enabled**, **show interface expand detail** Supported platforms: Firepower 4100/9300
Firepower 1010 hardware switch support	6.5	Any	The Firepower 1010 supports setting each Ethernet interface to be a switch port or a firewall interface. New/Modified screens: • **Devices** > **Device Management** > **Interfaces** • **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** • **Devices** > **Device Management** > **Interfaces** > **Add VLAN Interface**
Firepower 1010 PoE+ support on Ethernet 1/7 and Ethernet 1/8	6.5	Any	The Firepower 1010 supports Power over Ethernet+ (PoE+) on Ethernet 1/7 and Ethernet 1/8 when they are configured as switch ports. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** > **PoE**

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
VLAN subinterfaces for use with container instances	6.3.0	Any	To provide flexible physical interface use, you can create VLAN subinterfaces in FXOS and also share interfaces between multiple instances. New/Modified Secure Firewall Management Center screens: **Devices** > **Device Management** > **Edit** icon **>** **Interfaces** tab New/Modified Secure Firewall Chassis Manager screens: **Interfaces** > **All Interfaces** > **Add New** drop-down menu **Subinterface** New/Modified FXOS commands: **create subinterface, set vlan, show interface,show subinterface** Supported platforms: Firepower 4100/9300
Data-sharing interfaces for container instances	6.3.0	Any	To provide flexible physical interface use, you can share interfaces between multiple instances. New/Modified Secure Firewall Chassis Manager screens: **Interfaces** > **All Interfaces** > **Type** New/Modified FXOS commands: **set port-type data-sharing, show interface** Supported platforms: Firepower 4100/9300

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Integrated Routing and Bridging	6.2.0	Any	Integrated Routing and Bridging provides the ability to route between a bridge group and a routed interface. A bridge group is a group of interfaces that the Firewall Threat Defense bridges instead of routes. The Firewall Threat Defense is not a true bridge in that the Firewall Threat Defense continues to act as a firewall: access control between interfaces is controlled, and all of the usual firewall checks are in place. Previously, you could only configure bridge groups in transparent firewall mode, where you cannot route between bridge groups. This feature lets you configure bridge groups in routed firewall mode, and to route between bridge groups and between a bridge group and a routed interface. The bridge group participates in routing by using a Bridge Virtual Interface (BVI) to act as a gateway for the bridge group. Integrated Routing and Bridging provides an alternative to using an external Layer 2 switch if you have extra interfaces on the Firewall Threat Defense to assign to the bridge group. In routed mode, the BVI can be a named interface and can participate separately from member interfaces in some features, such as access rules and DHCP server. The following features that are supported in transparent mode are not supported in routed mode: clustering. The following features are also not supported on BVIs: dynamic routing and multicast routing. New/Modified screens: • **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** • **Devices** > **Device Management** > **Interfaces** > **Add Interfaces** > **Bridge Group Interface** Supported platforms: All except for the Firepower 2100 and the Firewall Threat Defense Virtual

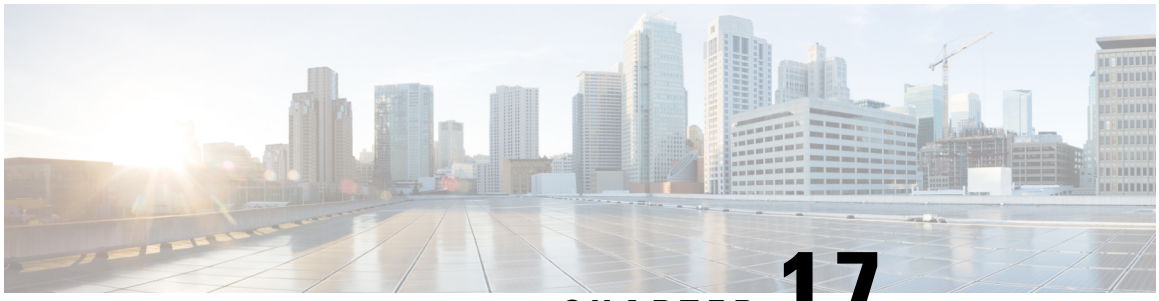

CHAPTER 17

Inline Sets and Passive Interfaces

You can configure IPS-only passive interfaces, passive ERSPAN interfaces, and inline sets.

- About IPS Interfaces, on page 893
- Requirements and Prerequisites for Inline Sets, on page 897
- Guidelines for Inline Sets and Passive Interfaces, on page 898
- Configure a Passive Interface, on page 900
- Configure an Inline Set, on page 901
- History for Inline Sets and Passive Interfaces, on page 905

About IPS Interfaces

IPS interfaces include passive interfaces, passive ERSPAN interfaces, and inline sets. IPS-only mode interfaces bypass many firewall checks and only support IPS security policy (Snort). You might want to implement IPS-only interfaces if you have a separate firewall protecting these interfaces and do not want the overhead of firewall functions.

IPS-only mode interfaces bypass many firewall checks and only support IPS security policy (Snort). You might want to implement IPS-only interfaces if you have a separate firewall protecting these interfaces and do not want the overhead of firewall functions.

Note The firewall mode only affects regular firewall interfaces, and not IPS-only interfaces such as inline sets or passive interfaces. IPS-only interfaces can be used in both firewall modes.

Inline Sets

An inline set acts like a bump on the wire, and binds one or more interface pairs together to slot into an existing network. This function allows the Firewall Threat Defense to be installed in any network environment without the configuration of adjacent network devices. Inline interfaces receive all traffic unconditionally, but all traffic received on these interfaces is retransmitted out of the other interface in the inline pair unless explicitly dropped. When you have multiple inline pairs in an inline set, traffic can only pass between the interfaces in the pair; it can't pass between interfaces in different pairs.

With tap mode, the Firewall Threat Defense is deployed inline, but the network traffic flow is undisturbed. Instead, the Firewall Threat Defense makes a copy of each packet so that it can analyze the packets. Note that

rules of these types do generate intrusion events when they are triggered, and the table view of intrusion events indicates that the triggering packets would have dropped in an inline deployment. There are benefits to using tap mode with FTDs that are deployed inline. For example, you can set up the cabling between the Firewall Threat Defense and the network as if the Firewall Threat Defense were inline and analyze the kinds of intrusion events the Firewall Threat Defense generates. Based on the results, you can modify your intrusion policy and add the drop rules that best protect your network without impacting its efficiency. When you are ready to deploy the Firewall Threat Defense inline, you can disable tap mode and begin dropping suspicious traffic without having to reconfigure the cabling between the Firewall Threat Defense and the network.

Note Tap mode *significantly* impacts Firewall Threat Defense performance, depending on the traffic.

Note Inline sets might be familiar to you as "transparent inline sets," but the inline interface type is unrelated to the transparent firewall mode or the firewall-type interfaces.

Multiple Inline Pairs and Asynchronous Routing

You can configure the interfaces to route traffic between a host on your network and external hosts through different inline pairs, depending on whether the traffic is inbound or outbound. This is an *asynchronous routing* configuration. If you deploy asynchronous routing, but you include only one inline pair in an inline set, the device might not correctly analyze your network traffic because it might see only half of the traffic.

Adding multiple inline pairs to the same inline set lets the system identify the inbound and outbound traffic as part of the same traffic flow. For passive interfaces only, you can also achieve this by including the interface pairs in the same security zone.

Figure 383: Asynchronous Routing

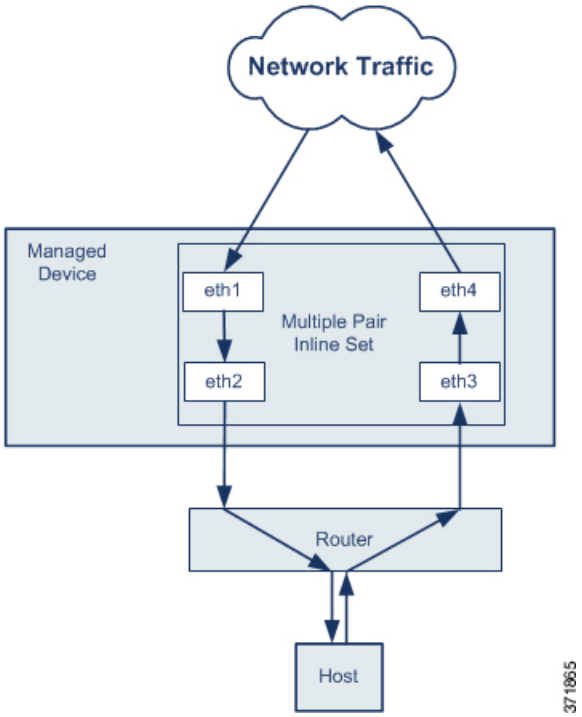

Note If you assign multiple inline pairs to a single inline set, but you experience issues with duplicate traffic, you might need to reassign your inline pairs to separate inline sets or modify your security zones.

If fragments of a packet are received on different interface pairs, they are not reassembled and get dropped. Make sure all the fragments of the packet are received and sent on the same interface pair.

Passive Interfaces

Passive interfaces monitor traffic flowing across a network using a switch SPAN or mirror port. The SPAN or mirror port allows for traffic to be copied from other ports on the switch. This function provides the system visibility within the network without being in the flow of network traffic. When you configure the Firewall Threat Defense in a passive deployment, the Firewall Threat Defense cannot take certain actions such as blocking or shaping traffic. Passive interfaces receive all traffic unconditionally. and no traffic received on these interfaces is retransmitted. Encapsulated remote switched port analyzer (ERSPAN) interfaces allow you to monitor traffic from source ports distributed over multiple switches, and uses GRE to encapsulate the traffic. ERSPAN interfaces are only allowed when the Firewall Threat Defense is in routed firewall mode.

Note Using SR-IOV interfaces as passive interfaces on NGFWv is not supported on some Intel network adapters (such as Intel X710 or 82599) using SR-IOV drivers due to a promiscuous mode restriction. In such cases, use a network adapter that supports this functionality. See Intel Ethernet Products for more information on Intel network adapters.

About Hardware Bypass for Inline Sets

For certain interface modules on the supported models (see Requirements and Prerequisites for Inline Sets, on page 897), you can enable the Hardware Bypass feature. Hardware Bypass ensures that traffic continues to flow between an inline interface pair during a power outage. This feature can be used to maintain network connectivity in the case of software or hardware failures.

Hardware Bypass Triggers

Hardware Bypass can be triggered in the following scenarios:

- Firewall Threat Defense crash
- Firewall Threat Defense reboot
- Security Module reboot
- Chassis crash
- Chassis reboot
- Manual trigger
- Chassis power loss
- Security Module power loss

Note Hardware bypass is intended for unplanned/unexpected failure scenarios, and is not automatically triggered during planned software upgrades. Hardware bypass only engages at the end of a planned upgrade process, when the Firewall Threat Defense application reboots.

Hardware Bypass Switchover

When switching from normal operation to hardware bypass or from hardware bypass back to normal operation, traffic may be interrupted for several seconds. A number of factors can affect the length of the interruption; for example, copper port auto-negotiation; behavior of the optical link partner such as how it handles link faults and de-bounce timing; spanning tree protocol convergence; dynamic routing protocol convergence; and so on. During this time, you may experience dropped connections.

You may also experience dropped connections due to application identification errors when analyzing connections midstream after the return to normal operations.

Snort Fail Open vs. Hardware Bypass

For inline sets other than those in tap mode, you can use the Snort Fail Open option to either drop traffic or allow traffic to pass without inspection when the Snort process is busy or down. Snort Fail Open is supported on all inline sets except those in tap mode, not just on interfaces that support Hardware Bypass.

The Hardware Bypass functionality allows traffic to flow during a hardware failure, including a complete power outage, and certain limited software failures. A software failure that triggers Snort Fail Open does not trigger a Hardware Bypass.

Hardware Bypass Status

If the system has power, then the Bypass LED indicates the Hardware Bypass status. See the Firepower chassis hardware installation guide for LED descriptions.

Requirements and Prerequisites for Inline Sets

User Roles

- Admin
- Access Admin
- Network Admin

Hardware Bypass Support

The Firewall Threat Defense supports Hardware Bypass for interface pairs on specific network modules on the following models:

- Secure Firewall 3100
- Firepower 4100
- Secure Firewall 4200
- Firepower 9300

Note The ISA 3000 has a separate implementation for Hardware Bypass, which you can enable using FlexConfig only (see FlexConfig Policies, on page 2667). Do not use this chapter to configure ISA 3000 Hardware Bypass.

Note You can use Hardware Bypass interfaces as regular interfaces without the Hardware Bypass feature enabled.

The supported Hardware Bypass network modules for these models include:

- Secure Firewall 3100:
  - 6-port 1G SFP Fail-to-Wire Network Module, SX (multimode) (FPR3K-XNM-6X1SXF)
  - 6-port 10G SFP Fail-to-Wire Network Module, SR (multimode) (FPR3K-XNM-6X10SRF)
  - 6-port 10G SFP Fail-to-Wire Network Module, LR (single mode) (FPR3K-XNM-6X10LRF)
  - 6-port 25G SFP Fail-to-Wire Network Module, SR (multimode) (FPR3K-XNM-X25SRF)
  - 6-port 25G Fail-to-Wire Network Module, LR (single mode) (FPR3K-XNM-6X25LRF)
  - 8-port 1G Copper Fail-to-Wire Network Module, RJ45 (copper) (FPR3K-XNM-8X1GF)
- Secure Firewall 4200:

- 6-port 1G SFP Fail-to-Wire Network Module, SX (multimode) (FPR4K-XNM-6X1SXF)
- 6-port 10G SFP Fail-to-Wire Network Module, SR (multimode) (FPR4K-XNM-6X10SRF)
- 6-port 10G SFP Fail-to-Wire Network Module, LR (single mode) (FPR4K-XNM-6X10LRF)
- 6-port 25G SFP Fail-to-Wire Network Module, SR (multimode) (FPR4K-XNM-X25SRF)
- 6-port 25G Fail-to-Wire Network Module, LR (single mode) (FPR4K-XNM-6X25LRF)
- 8-port 1G Copper Fail-to-Wire Network Module, RJ45 (copper) (FPR4K-XNM-8X1GF)

- Firepower 4100:
  - Firepower 6-port 1G SX FTW Network Module single-wide (FPR4K-NM-6X1SX-F)
  - Firepower 6-port 10G SR FTW Network Module single-wide (FPR4K-NM-6X10SR-F)
  - Firepower 6-port 10G LR FTW Network Module single-wide (FPR4K-NM-6X10LR-F)
  - Firepower 2-port 40G SR FTW Network Module single-wide (FPR4K-NM-2X40G-F)
  - Firepower 8-port 1G Copper FTW Network Module single-wide (FPR-NM-8X1G-F)

- Firepower 9300:
  - Firepower 6-port 10G SR FTW Network Module single-wide (FPR9K-NM-6X10SR-F)
  - Firepower 6-port 10G LR FTW Network Module single-wide (FPR9K-NM-6X10LR-F)
  - Firepower 2-port 40G SR FTW Network Module single-wide (FPR9K-NM-2X40G-F)

Hardware Bypass can only use the following port pairs:

- 1 & 2
- 3 & 4
- 5 & 6
- 7 & 8

Guidelines for Inline Sets and Passive Interfaces

Firewall Mode

- ERSPAN interfaces are only allowed when the device is in routed firewall mode.

Clustering

- Link State Propagation for an inline set is not supported with clustering.
- IPS-only interfaces are not supported in Individual interface mode.

Multi-Instance Mode

- Multi-instance shared interfaces are not supported. You must use an unshared interface.
- Multi-instance chassis-defined subinterfaces are not supported. You must use a physical interface or EtherChannel.

General Guidelines

- Inline sets and passive interfaces support physical interfaces and EtherChannels only, and cannot use VLANs or other virtual interfaces, including multi-instance chassis-defined subinterfaces.
- Since IPS interfaces do not support regular firewall protections, the IPS security policy (Snort) requires that traffic flows pass through the same Firewall Threat Defense so that all traffic can be inspected.
- Bidirectional Forwarding Detection (BFD) echo packets are not allowed through the Firewall Threat Defense when using inline sets. If there are two neighbors on either side of the Firewall Threat Defense running BFD, then the Firewall Threat Defense will drop BFD echo packets because they have the same source and destination IP address and appear to be part of a LAND attack.
- For inline sets and passive interfaces, the Firewall Threat Defense supports up to two 802.1Q headers in a packet (also known as Q-in-Q support), with the exception of the Firepower 4100/9300, which only supports one 802.1Q header. **Note:** Firewall-type interfaces do not support Q-in-Q, and only support one 802.1Q header.

Hardware Bypass Guidelines

- Hardware Bypass ports are supported only for inline sets.
- Hardware Bypass ports cannot be part of an EtherChannel.
- Hardware Bypass is not supported in high availability mode.
- Hardware Bypass ports are supported with intra-chassis clustering on the Firepower 9300. Ports are placed in Hardware Bypass mode when the last unit in the chassis fails. Inter-chassis clustering is not supported, because inter-chassis clustering only supports Spanned EtherChannels; Hardware Bypass ports cannot be part of an EtherChannel.
- If all modules in an intra-chassis cluster on the Firepower 9300 fail, then Hardware Bypass is triggered on the final unit, and traffic continues to pass. When units come back up, Hardware Bypass returns to standby mode. However, when you use rules that match application traffic, those connections may be dropped and need to be reestablished. Connections are dropped because state information is not retained on the cluster unit, and the unit cannot identify the traffic as belonging to an allowed application. To avoid a traffic drop, use a port-based rule instead of an application-based rule, if appropriate for your deployment.
- You can use Hardware Bypass interfaces as regular interfaces without the Hardware Bypass feature enabled.

Unsupported Firewall Features on IPS Interfaces

- DHCP server
- DHCP relay

- DHCP client
- TCP Intercept
- Routing
- NAT
- VPN
- Application inspection
- QoS
- NetFlow
- VXLAN

Configure a Passive Interface

This section describes how to:

- Enable the interface. By default, interfaces are disabled.
- Set the interface mode to Passive or ERSPAN. For ERSPAN interfaces, you will set the ERSPAN parameters and the IP address.
- Change the MTU. By default, the MTU is set to 1500 bytes. For more information about the MTU, see About the MTU, on page 877.
- Set a specific speed and duplex (if available). By default, speed and duplex are set to Auto.

Note For the Secure Firewall Threat Defense on the FXOS chassis, you configure basic interface settings on the Firepower 4100/9300. See Configure a Physical Interface, on page 332 for more information.

Before you begin

- If you are using EtherChannels, add them according to Configure an EtherChannel, on page 786.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 In the **Mode** drop-down list, choose **Passive** or **Erspan**.

Step 4 Enable the interface by checking the **Enabled** check box.

Step 5 In the **Name** field, enter a name up to 48 characters in length.

Step 6 From the **Security Zone** drop-down list, choose a security zone or add a new one by clicking **New**.

Step 7 (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

Step 8 (Optional) On **General**, set the **MTU** between 64 and 9198 bytes; for the Secure Firewall Threat Defense Virtual and Secure Firewall Threat Defense on the FXOS chassis, the maximum is 9000 bytes.

The default is 1500 bytes.

Step 9 For ERSPAN interfaces, set the following parameters:

- **Flow Id**—Configure the ID used by the source and destination sessions to identify the ERSPAN traffic, between 1 and 1023. This ID must also be entered in the ERSPAN destination session configuration.
- **Source IP**—Configure the IP address used as the source of the ERSPAN traffic.

Step 10 For ERSPAN interfaces, set the IPv4 address and mask on **IPv4**.

Step 11 (Optional) Set the duplex and speed by clicking **Hardware Configuration**.

The exact speed and duplex options depend on your hardware.

- **Duplex**—Choose **Full**, **Half**, or **Auto**. Auto is the default.
- **Speed**—Choose **10**, **100**, **1000**, or **Auto**. Auto is the default.

Step 12 Click **OK**.

Step 13 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure an Inline Set

This section enables and names two physical interfaces or EtherChannels per inline pair that you can add to an inline set. You can add multiple inline pairs per inline set. You can also optionally enable Hardware Bypass for supported inline pairs.

Note For the Firepower 4100/9300, you configure basic interface settings in FXOS on the chassis. See Configure a Physical Interface, on page 332 for more information.

Before you begin

- If you are using EtherChannels, add them according to Configure an EtherChannel, on page 786.
- We recommend that you set STP PortFast for STP-enabled switches that connect to the Firewall Threat Defense inline pair interfaces. This setting is especially useful for Hardware Bypass configurations and can reduce bypass times.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 For each interface in the inline pair, name the interface and enable it. You can also configure other hardware settings. Be sure to match the hardware settings for each interface in the inline pair. You can configure multiple pairs of interfaces.

a) Click **Edit** (✎) for the interface you want to edit.
b) In the **Name** field, enter a name up to 48 characters in length.
c) Enable the interface by checking the **Enabled** check box.
d) (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

e) Leave the **Mode** as **None**.

After you add this interface to an inline set, this field will show Inline for the mode.

f) Do not set the security zone yet; you must set it after you create the inline set later in this procedure.
g) (Optional) Set the duplex and speed by clicking **Hardware Configuration**.

The exact speed and duplex options depend on your hardware.

- **Duplex**—Choose **Full**, **Half**, or **Auto**. Auto is the default.
- **Speed**—Choose **10**, **100**, **1000**, or **Auto**. Auto is the default.

h) Click **OK**.

Do not set any other settings for this interface.

Step 3 Click **Inline Sets**.

Step 4 Click **Add Inline Set**.

Figure 384: Add Inline Set

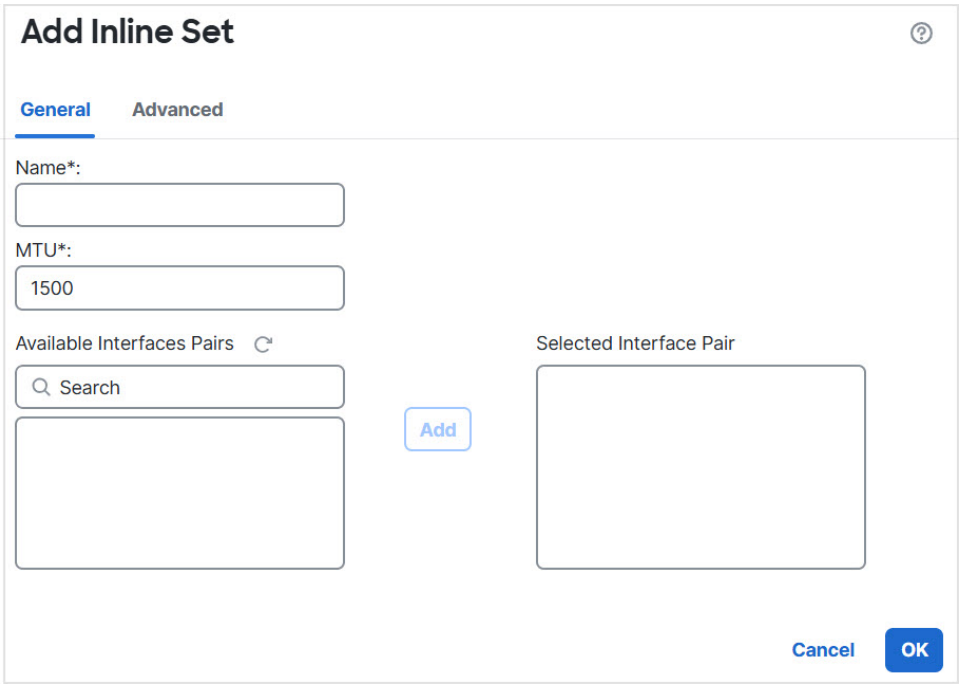

The **Add Inline Set** dialog box appears with **General** selected.

Step 5 In the **Name** field, enter a name for the set.

Step 6 (Optional) Change the **MTU** to enable jumbo frames.

For inline sets, the MTU setting is not used. However, the jumbo frame setting *is* relevant to inline sets; jumbo frames enable the inline interfaces to receive packets up to 9000 bytes. To enable jumbo frames, you must set the MTU of *any* interface on the device above 1500 bytes.

Step 7 Configure Hardware Bypass.

a) For the **Bypass** mode, choose one of the following options:

- **Disabled**—Set Hardware Bypass to disabled for interfaces where Hardware Bypass is supported, or use interfaces where Hardware Bypass is not supported.
- **Standby**—Set Hardware Bypass to the standby state on supported interfaces. Only pairs of Hardware Bypass interfaces are shown. In the standby state, the interfaces remain in normal operation until there is a trigger event.
- **Bypass-Force**—Manually forces the interface pair to go into a bypass state. **Inline Sets** shows **Yes** for any interface pairs that are in Bypass-Force mode.

b) In the **Available Interfaces Pairs** area, click a pair and then click **Add** to move it to the **Selected Interface Pair** area.

All possible pairings between named and enabled interfaces with the mode set to None show in this area.

Step 8 (Optional) Click **Advanced** to set the following optional parameters:

- **Tap Mode**—Set to inline tap mode.

Note that you cannot enable this option and strict TCP enforcement on the same inline set.

Note
If you need to enable or disable the Tap mode, you should do so during a maintenance window. Changing the mode while the device is passing traffic can cause traffic disruption.

Note
Tap mode *significantly* impacts the Firewall Threat Defense performance, depending on the traffic.

- **Propagate Link State**—Configure link state propagation.

  Link state propagation automatically brings down the second interface in the inline interface pair when one of the interfaces in an inline set goes down. When the downed interface comes back up, the second interface automatically comes back up, also. In other words, if the link state of one interface changes, the device senses the change and updates the link state of the other interface to match it. Note that devices require up to 4 seconds to propagate link state changes. Link state propagation is especially useful in resilient network environments where routers are configured to reroute traffic automatically around network devices that are in a failure state.

  Note
  Do not enable **Propagate Link State** when using clustering.

- **Snort Fail Open**—Enable or disable either or both of the **Busy** and **Down** options if you want new and existing traffic to pass without inspection (enabled) or drop (disabled) when the Snort process is busy or down.

  By default, traffic passes without inspection when the Snort process is down, and drops when it is busy.

  When the Snort process is:

  - Busy—It cannot process traffic fast enough because traffic buffers are full, indicating that there is more traffic than the device can handle, or because of other software resource issues.
  - Down—It is restarting because you deployed a configuration that requires it to restart. See Configurations that Restart the Snort Process When Deployed or Activated, on page 256.

  When the Snort process is down and comes back up, it inspects *new* connections. To prevent false positives and false negatives, it does not inspect existing connections on inline, routed, or transparent interfaces because initial session information might have been lost while it was down.

  Note
  When Snort fails open, features that rely on the Snort process do not function. These include application control and deep inspection. The system performs only basic access control using simple, easily determined transport and network layer characteristics.

  Note
  The **Strict TCP Enforcement** option is not supported.

Step 9 Set the security zone for each interface.

a) Click **Interfaces**.
b) Click **Edit** (✎) for the member interface.
c) From the **Security Zone** drop-down list, choose a security zone or add a new one by clicking **New**.

You can only set the zone after you add the interface to the inline set; adding it to an inline set configures the mode to Inline and lets you choose inline-type security zones.

d) Click **OK**.

Step 10 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

History for Inline Sets and Passive Interfaces

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Hardware bypass support on the Secure Firewall 4200 for supported network modules	7.4.0	7.4.0	The Secure Firewall 4200 supports hardware bypass functionality when using the hardware bypass network modules. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** Supported platforms: Secure Firewall 4200
Hardware bypass support on the Secure Firewall 3100 for supported network modules	7.2	Any	The Secure Firewall 3100 now supports hardware bypass functionality when using the hardware bypass network modules. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** Supported platforms: Secure Firewall 3100

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Synchronization between the Firewall Threat Defense operational link state and the physical link state for the Firepower 4100/9300	6.7	Any	The Firepower 4100/9300 chassis can now synchronize the Firewall Threat Defense operational link state with the physical link state for data interfaces. Currently, interfaces will be in an Up state as long as the FXOS admin state is up and the physical link state is up. The Firewall Threat Defense application interface admin state is not considered. Without synchronization from Firewall Threat Defense, data interfaces can be in an Up state physically before the Firewall Threat Defense application has completely come online, for example, or can stay Up for a period of time after you initiate a shutdown. For inline sets, this state mismatch can result in dropped packets because external routers may start sending traffic to the Firewall Threat Defense before the Firewall Threat Defense can handle it. This feature is disabled by default, and can be enabled per logical device in FXOS. **Note** This feature is not supported for clustering, container instances, or Firewall Threat Defense with a Radware vDP decorator. It is also not supported for ASA. New/Modified Firepower Chassis Manager screens: **Logical Devices** > **Enable Link State** New/Modified FXOS commands: **set link-state-sync enabled**, **show interface expand detail** Supported platforms: Firepower 4100/9300
Hardware bypass support on the Firepower 2130 and 2140 for supported network modules	6.3.0	Any	The Firepower 2130 and 2140 now support hardware bypass functionality when using the hardware bypass network modules. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** Supported platforms: Firepower 2130 and 2140
Support for EtherChannels in Firewall Threat Defense inline sets or passive interface	6.2.0	Any	You can now use EtherChannels in a Firewall Threat Defense inline set or passive interface.
Hardware bypass support on the Firepower 4100/9300 for supported network modules	6.1.0	Any	Hardware Bypass ensures that traffic continues to flow between an inline interface pair during a power outage. This feature can be used to maintain network connectivity in the case of software or hardware failures. New/Modified screens: **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** Supported platforms: Firepower 4100/9300

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Inline set link state propagation support for the Firewall Threat Defense	6.1.0	Any	When you configure an inline set in the Firewall Threat Defense application and enable link state propagation, the Firewall Threat Defense sends inline set membership to the FXOS chassis. Link state propagation means that the chassis automatically brings down the second interface in the inline interface pair when one of the interfaces in an inline set goes down. New/Modified FXOS commands: **show fault |grep link-down**, **show interface detail**

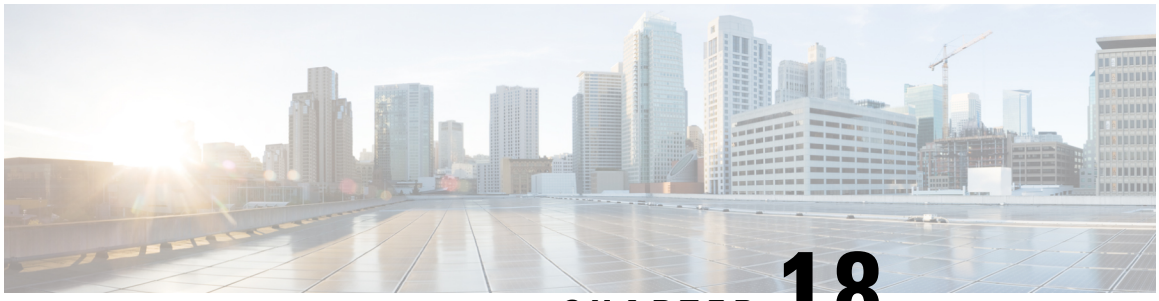

CHAPTER 18

DHCP and DDNS

The following topics explain DHCP and DDNS services and how to configure them on Threat Defense devices.

- About DHCP and DDNS Services, on page 909
- Requirements and Prerequisites for DHCP and DDNS, on page 911
- Guidelines for DHCP and DDNS Services, on page 911
- Configure the DHCPv4 Server, on page 913
- Configure the DHCPv6 Stateless Server, on page 914
- Configure the DHCP Relay Agent, on page 919
- Configure Dynamic DNS, on page 920
- History for DHCP and DDNS, on page 926

About DHCP and DDNS Services

The following topics describe the DHCP server, DHCP relay agent, and DDNS update.

About the DHCPv4 Server

DHCP provides network configuration parameters, such as IP addresses, to DHCP clients. The Firewall Threat Defense device can provide a DHCP server to DHCP clients attached to Firewall Threat Defense device interfaces. The DHCP server provides network configuration parameters directly to DHCP clients.

An IPv4 DHCP client uses a broadcast rather than a multicast address to reach the server. The DHCP client listens for messages on UDP port 68; the DHCP server listens for messages on UDP port 67.

DHCP Options

DHCP provides a framework for passing configuration information to hosts on a TCP/IP network. The configuration parameters are carried in tagged items that are stored in the Options field of the DHCP message and the data are also called options. Vendor information is also stored in Options, and all of the vendor information extensions can be used as DHCP options.

For example, Cisco IP Phones download their configuration from a TFTP server. When a Cisco IP Phone starts, if it does not have both the IP address and TFTP server IP address preconfigured, it sends a request with option 150 or 66 to the DHCP server to obtain this information.

- DHCP option 150 provides the IP addresses of a list of TFTP servers.

- DHCP option 66 gives the IP address or the hostname of a single TFTP server.
- DHCP option 3 sets the default route.

A single request might include both options 150 and 66. In this case, the DHCP server provides values for both options in the response if they are already configured on the Firewall Threat Defense.

You can use advanced DHCP options to provide DNS, WINS, and domain name parameters to DHCP clients; DHCP option 15 is used for the DNS domain suffix. You can also use the DHCP automatic configuration setting to obtain these values or define them manually. When you use more than one method to define this information, it is passed to DHCP clients in the following sequence:

1. Manually configured settings.
2. Advanced DHCP options settings.
3. DHCP automatic configuration settings.

For example, you can manually define the domain name that you want the DHCP clients to receive and then enable DHCP automatic configuration. Although DHCP automatic configuration discovers the domain together with the DNS and WINS servers, the manually defined domain name is passed to DHCP clients with the discovered DNS and WINS server names, because the domain name discovered by the DHCP automatic configuration process is superseded by the manually defined domain name.

About the DHCPv6 Stateless Server

For clients that use StateLess Address Auto Configuration (SLAAC) in conjunction with the Prefix Delegation feature (Enable the IPv6 Prefix Delegation Client, on page 867), you can configure the Firewall Threat Defense to provide information such as the DNS server or domain name when they send Information Request (IR) packets to the Firewall Threat Defense by defining a DHCP IPv6 Pool and assigning it to the DHCPv6 server. The Firewall Threat Defense only accepts IR packets and does not assign addresses to the clients. You will configure the client to generate its own IPv6 address by enabling IPv6 autoconfiguration on the client. Enabling stateless autoconfiguration on a client configures IPv6 addresses based on prefixes received in Router Advertisement messages; in other words, based on the prefix that the Firewall Threat Defense received using Prefix Delegation.

About the DHCP Relay Agent

You can configure a DHCP relay agent to forward DHCP requests received on an interface to one or more DHCP servers. DHCP clients use UDP broadcasts to send their initial DHCPDISCOVER messages because they do not have information about the network to which they are attached. If the client is on a network segment that does not include a server, UDP broadcasts normally are not forwarded by the Firewall Threat Defense device because it does not forward broadcast traffic. The DHCP relay agent lets you configure the interface of the Firewall Threat Defense device that is receiving the broadcasts to forward DHCP requests to a DHCP server on another device.

Requirements and Prerequisites for DHCP and DDNS

Model Support

Threat Defense

User Roles

- Admin
- Access Admin
- Network Admin

Guidelines for DHCP and DDNS Services

This section includes guidelines and limitations that you should check before configuring DHCP and DDNS services.

Firewall Mode

- DHCP Relay is not supported in transparent firewall mode or in routed mode on the BVI or bridge group member interface.
- DHCP Server is supported in transparent firewall mode on a bridge group member interface. In routed mode, the DHCP server is supported on the BVI interface, not the bridge group member interface. The BVI must have a name for the DHCP server to operate.
- DDNS is not supported in transparent firewall mode or in routed mode on the BVI or bridge group member interface.
- DHCPv6 stateless server is not supported in transparent firewall mode or in routed mode on the BVI or bridge group member interface.

Clustering

- DHCPv6 stateless server is not supported with clustering.

IPv6

Supports IPv6 for DHCP stateless server and DHCP Relay.

DHCPv4 Server

- The maximum available DHCP pool is 256 addresses.
- You can configure a DHCP server on any interface with a name and IP address, such as a physical interface, a subinterface, or a BVI in routed mode.

- You can configure only one DHCP server on each interface. Each interface can have its own pool of addresses to use. However the other DHCP settings, such as DNS servers, domain name, options, ping timeout, and WINS servers, are configured globally and used by the DHCP server on all interfaces.
- You cannot configure an interface as a DHCP client if that interface also has DHCP server enabled; you must use a static IP address.
- You cannot configure both a DHCP server and DHCP relay on the same device, even if you want to enable them on different interfaces; you can only configure one type of service.
- The Firewall Threat Defense does not support DHCP clients behind DHCP relay servers; the client must be directly connected to the Firewall Threat Defense.
- The DHCP server does not support BOOTP requests.

DHCPv6 Server

The DHCPv6 Stateless server cannot be configured on an interface where the DHCPv6 address, Prefix Delegation client, or DHCPv6 relay is configured.

DHCP Relay

- You can configure a maximum of 10 DHCPv4 relay servers per virtual router, global (VRF) and interface-specific servers combined, with a maximum of 4 servers per interface.
- You can configure a maximum of 10 DHCPv6 relay servers per virtual router. Interface-specific servers for IPv6 are not supported.
- You cannot configure both a DHCP server and DHCP relay on the same device, even if you want to enable them on different interfaces; you can only configure one type of service.
- DHCP relay services are not available in transparent firewall mode or in routed mode on the BVI or bridge group member interface. You can, however, allow DHCP traffic through using an access rule. To allow DHCP requests and replies through the Firewall Threat Defense, you need to configure two access rules, one that allows DCHP requests from the inside interface to the outside (UDP destination port 67), and one that allows the replies from the server in the other direction (UDP destination port 68).
- For IPv4, clients must be directly-connected to the Firewall Threat Defense and cannot send requests through another relay agent or a router. For IPv6, the Firewall Threat Defense supports packets from another relay server.
- The DHCP clients must be on different interfaces from the DHCP servers to which the Firewall Threat Defense relays requests.
- You cannot enable DHCP Relay on an interface in a traffic zone.

DDNS Service

The firewall's DDNS supports only DynDNS service. Hence, ensure that the DDNS is configured with update URL in the following syntax:

https://*username***:***password***@***provider-domain***/***path***?hostname=<h>&myip=<a>**.

Configure the DHCPv4 Server

See the following steps to configure a DHCPv4 server.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select **DHCP > DHCP Server**.

Step 3 Configure the following DHCP server options:

- **Ping Timeout**—The amount of time in milliseconds that Firewall Threat Defense device waits to time out a DHCP ping attempt. Valid values range from 10 to 10000 milliseconds. The default value is 50 milliseconds.

  To avoid address conflicts, the Firewall Threat Defense device sends two ICMP ping packets to an address before assigning that address to a DHCP client.

- **Lease Length**—The amount of time in seconds that the client may use its allocated IP address before the lease expires. Valid values range from 300 to 1048575 seconds. The default value is 3600 seconds (1 hour).

- (Routed mode) **Auto-configuration**—Enables DHCP auto configuration on the Firewall Threat Defense device. Auto-configuration enables the DHCP server to provide the DHCP clients with the DNS server, domain name, and WINS server information obtained from a DHCP client running on the specified interface. Otherwise, you can disable auto configuration and add the values yourself in Step 4.

- (Routed mode) **Interface**—Specifies the interface to be used for auto configuration. For a device with virtual routing capability, this interface can only be a global virtual router interface.

Step 4 To override auto-configured settings, do the following:

- Enter the domain name of the interface. For example, your device may be in the Your_Company domain.

- From the drop-down list, choose the DNS servers (primary and secondary) configured for the interface. To add a new DNS server, see Creating Network Objects, on page 1403.

- From the drop-down list, choose the WINS servers (primary and secondary) configured for the interface. To add a new WINS server, see Creating Network Objects, on page 1403.

Step 5 Select **Server**, click **Add**, and configure the following options:

- **Interface**—Choose the interface from the drop-down list. In transparent mode, specify a named bridge group member interface. In routed mode, specify a named routed interface or a named BVI; do not specify the bridge group member interface. Note that each bridge group member interface for the BVI must also be named for the DHCP server to operate.

- **Address Pool**—The range of IP addresses from lowest to highest that is used by the DHCP server. The range of IP addresses must be on the same subnet as the selected interface and cannot include the IP address of the interface itself.

- **Enable DHCP Server**—Enables the DHCP server on the selected interface.

Step 6 Click **OK** to save the DHCP server configuration.

Step 7 (Optional) Select **Advanced**, click **Add**, and specify the type of information you want the option to return to the DHCP client:

- **Option Code**—The Firewall Threat Defense device supports the DHCP options listed in RFC 2132, RFC 2562, and RFC 5510 to send information. All DHCP options (1 through 255) are supported except for 1, 12, 50–54, 58–59, 61, 67, and 82. See About the DHCPv4 Server, on page 909for more information on DHCP option codes.

  Note
  The Firewall Threat Defense device does not verify that the option type and value that you provide match the expected type and value for the option code, as defined in RFC 2132. For more information about option codes and their associated types and expected values, see RFC 2132.

- **Type**—DHCP option type. Available options include **IP**, **ASCII**, and **HEX**. If you chose IP, you must add IP addresses in the IP Address fields. If you chose ASCII, you must add the ASCII value in the ASCII field. If you chose HEX, you must add the HEX value in the HEX field.
- **IP Address 1** and **IP Address 2**—The IP address(es) to be returned with this option code. To add a new IP address, see Creating Network Objects, on page 1403.
- **ASCII**—The ASCII value that is returned to the DHCP client. The string cannot include spaces.
- **HEX**—The HEX value that is returned to the DHCP client. The string must have an even number of digits and no spaces. You do not need to use a 0x prefix.

Step 8 Click **OK** to save the option code configuration.

Step 9 Click **Save** on the DHCP page to save your changes.

Step 10 To view DHCP bindings, use the following command.

show dhcpd binding

Example:

```
> show dhcpd binding
IP Address Client-id         Lease Expiration  Type
10.0.1.100 0100.a0c9.868e.43 84985 seconds     automatic
```

Configure the DHCPv6 Stateless Server

For clients that use StateLess Address Auto Configuration (SLAAC) in conjunction with the Prefix Delegation feature, you can configure the Firewall Threat Defense to provide information such as the DNS server or domain name when they send Information Request (IR) packets to the Firewall Threat Defense.

Create the DHCP IPv6 Pool

Create a DHCP IPv6 Pool for use with the DHCPv6 server. The DHCPv6 server provides information such as the DNS server or domain name when they clients send Information Request (IR) packets to the Firewall Threat Defense. The DHCP IPv6 Pool defines the parameters to sent in IR messages.

This feature is only supported in routed mode. This feature is not supported in clustering or High Availability.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **DHCP IPv6 Pool** from the list of object types.

Step 3 Click **Add** ( ).

Step 4 Configure the **DNS Server** and **Domain Name**.

You can either manually define the values and click **Add**, or you can check **Import** to use one or more parameters that the Firewall Threat Defense obtained from the DHCPv6 server on the Prefix Delegation client interface. You can mix and match manually-configured parameters with imported parameters; however, you cannot configure the same parameter manually and also use **Import**.

Figure 385: Manually Define Values

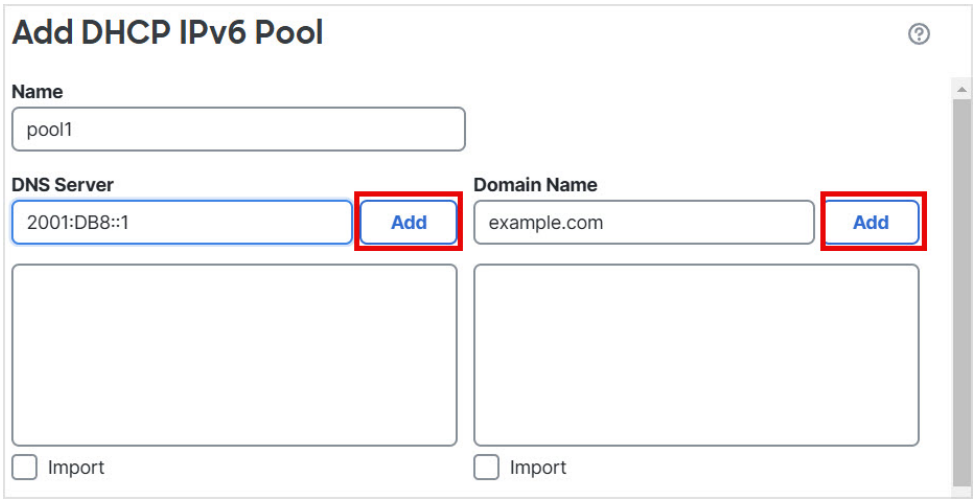

Figure 386: Import Values

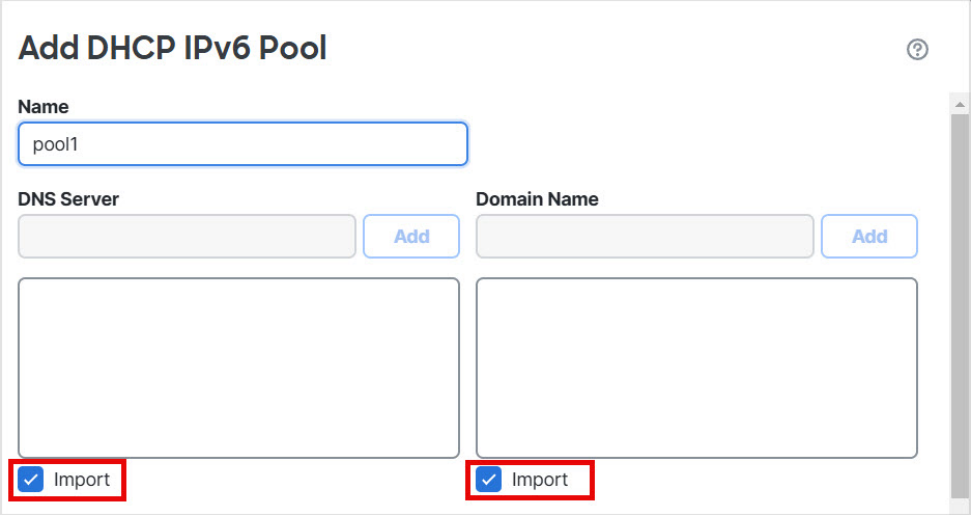

Step 5 Define **Other Server Options**.

You can define the domain name and IP address for the following servers:

- NIS
- NISP
- SIP
- SNTP

a) Click **Add** ( ).

Figure 387: Other Server Options

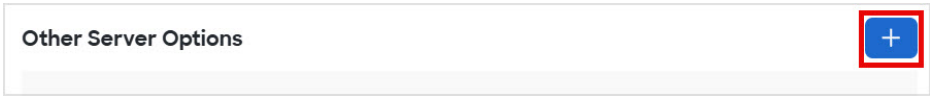

b) Choose the server type under **Option**, and either manually define the **Domain Name** and **Address**, or check **Import**.

Figure 388: Define Server Domain Name and Address

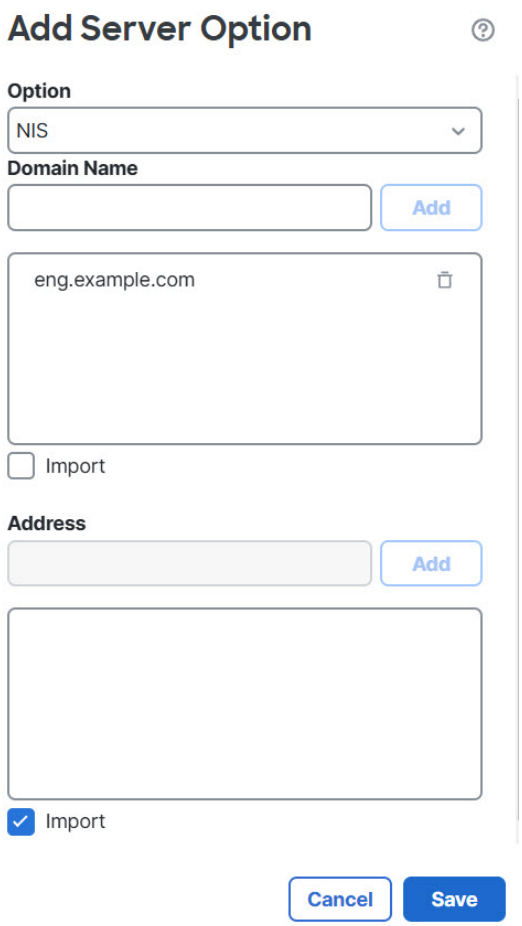

Import uses one or more parameters that the Firewall Threat Defense obtained from the DHCPv6 server on the Prefix Delegation client interface. You can mix and match manually-configured parameters with imported parameters; however, you cannot configure the same parameter manually and also use **Import**.

c) Click **Save**.
d) Repeat for each server type.

Step 6 Click **Save**.

Step 7 Use this pool with the DHCPv6 server. See Enable the DHCPv6 Stateless Server, on page 917.

Enable the DHCPv6 Stateless Server

For clients that use StateLess Address Auto Configuration (SLAAC) in conjunction with the Prefix Delegation feature (Enable the IPv6 Prefix Delegation Client, on page 867), you can configure the Firewall Threat Defense to provide information such as the DNS server or domain name when they send Information Request (IR) packets to the Firewall Threat Defense. The Firewall Threat Defense only accepts IR packets and does not assign addresses to the clients. You will configure the client to generate its own IPv6 address by enabling IPv6 autoconfiguration on the client. Enabling stateless autoconfiguration on a client configures IPv6 addresses

based on prefixes received in Router Advertisement messages; in other words, based on the prefix that the Firewall Threat Defense received using Prefix Delegation.

This feature is only supported in routed mode. This feature is not supported in clustering or High Availability.

Before you begin

Add a DHCP IPv6 Pool object. See Create the DHCP IPv6 Pool, on page 914. This object defines the server parameters included in the IR messages.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click the **IPv6** page, and then click **DHCP**.

Step 4 Click **DHCP Server Pool**, and choose the object you created earlier.

Figure 389: Enable the DHCPv6 Server

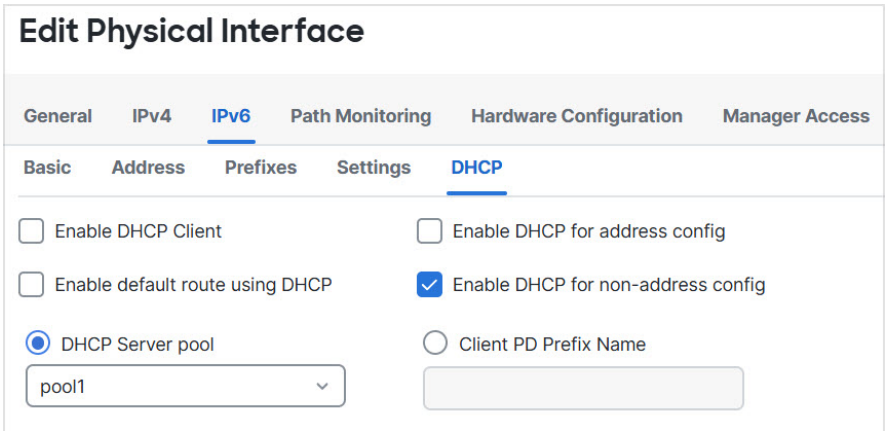

Step 5 Check **Enable DHCP for non-address config** to inform SLAAC clients about the DHCPv6 server.

This flag informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain additional information from DHCPv6, such as the DNS server address.

Step 6 Click **OK**.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure the DHCP Relay Agent

You can configure a DHCP relay agent to forward DHCP requests received on an interface to one or more DHCP servers. DHCP clients use UDP broadcasts to send their initial DHCPDISCOVER messages because they do not have information about the network to which they are attached. If the client is on a network segment that does not include a server, UDP broadcasts normally are not forwarded by the Firewall Threat Defense device because it does not forward broadcast traffic.

You can remedy this situation by configuring the interface of the Firewall Threat Defense device that is receiving the broadcasts to forward DHCP requests to a DHCP server on another interface.

Note DHCP Relay is not supported in transparent firewall mode.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select **DHCP** > **DHCP Relay**.

Step 3 In the **IPv4 Relay Timeout** and **IPv6 Relay Timeout** fields, enter the amount of time in seconds that the Firewall Threat Defense device waits to time out the DHCP relay agent. Valid values range from 1 to 3600 seconds. The default value is 60 seconds.

The timeout is for address negotiation through the local DHCP Relay agent.

Step 4 (Optional) Check **Trust All Information** to set all client interfaces as trusted.

You can configure interfaces as trusted interfaces to preserve DHCP Option 82. DHCP Option 82 is used by downstream switches and routers for DHCP snooping and IP Source Guard. Normally, if the Firewall Threat Defense DHCP relay agent receives a DHCP packet with Option 82 already set, but the giaddr field (which specifies the DHCP relay agent address that is set by the relay agent before it forwards the packet to the server) is set to 0, then the Firewall Threat Defense will drop that packet by default. You can preserve Option 82 and forward the packet by identifying an interface as a trusted interface.

Step 5 On **DHCP Relay Agent**, click **Add**, and configure the following options:

- **Interface**—The interface connected to the DHCP clients.
- **Enable IPv4 Relay**—Enables IPv4 DHCP Relay for this interface.
- **Set Route**—(For IPv4) Changes the default gateway address in the DHCP message from the server to that of the Firewall Threat Defense device interface that is closest to the DHCP client, which relayed the original DHCP request. This action allows the client to set its default route to point to the Firewall Threat Defense device even if the DHCP server specifies a different router. If there is no default router option in the packet, the Firewall Threat Defense device adds one containing the interface address.
- **Enable IPv6 Relay**—Enables IPv6 DHCP Relay for this interface.

Step 6 Click **OK** to save the DHCP relay agent changes.

Step 7 On **DHCP Servers**, click **Add**, and configure the following options:

Add the IPv4 and IPv6 server addresses as separate entries, even if they belong to the same server.

- **Server**—The IP address of the DHCP server. Chose an IP address from the drop-down list. To add a new one, see Creating Network Objects, on page 1403
- **Interface**—The interface to which the specified DHCP server is attached. The DHCP Relay agent and the DHCP server cannot be configured on the same interface.

Step 8 Click **OK** to save the DHCP server changes.

Step 9 Click **Save** on the DHCP page to save your changes.

Configure Dynamic DNS

When an interface uses DHCP IP addressing, the assigned IP address can change when the DHCP lease is renewed. When the interface needs to be reachable using a fully qualified domain name (FQDN), the IP address change can cause the DNS server resource records (RRs) to become stale. Dynamic DNS (DDNS) provides a mechanism to update DNS RRs whenever the IP address or hostname changes. You can also use DDNS for static or PPPoE IP addressing.

DDNS updates the following RRs on the DNS server: the A RR includes the name-to-IP address mapping, while the PTR RR maps addresses to names.

The Firewall Threat Defense supports the following DDNS update methods:

- Standard DDNS—The standard DDNS update method is defined by RFC 2136.

  With this method, the Firewall Threat Defense and the DHCP server use DNS requests to update the DNS RRs. The Firewall Threat Defense or DHCP server sends a DNS request to its local DNS server for information about the hostname and, based on the response, determines the main DNS server that owns the RRs. The Firewall Threat Defense or DHCP server then sends an update request directly to the main DNS server. See the following typical scenarios.

  - The Firewall Threat Defense updates the A RR, and the DHCP server updates the PTR RR.

    Typically, the Firewall Threat Defense "owns" the A RR, while the DHCP server "owns" the PTR RR, so both entities need to request updates separately. When the IP address or hostname changes, the Firewall Threat Defense sends a DHCP request (including the FQDN option) to the DHCP server to inform it that it needs to request a PTR RR update.

  - The DHCP server updates both the A and PTR RR.

    Use this scenario if the Firewall Threat Defense does not have the authority to update the A RR. When the IP address or hostname changes, the Firewall Threat Defense sends a DHCP request (including the FQDN option) to the DHCP server to inform it that it needs to request an A and PTR RR update.

  You can configure different ownership depending on your security needs and the requirements of the main DNS server. For example, for a static address, the Firewall Threat Defense should own the updates for both records.

- Web—The Web update method uses the DynDNS Remote API specification (https://help.dyn.com/remote-access-api/).

With this method when the IP address or hostname changes, the Firewall Threat Defense sends an HTTP request directly to a DNS provider with which you have an account.

Note For devices that were registered using zero-touch provisioning from the outside interface, DDNS is automatically enabled using the "fmcOnly" method, which is similar to the Web method. This method is only available for zero-touch provisioning devices. You can edit some options for this method using this screen, or delete the method and configure a different one. For more information about zero-touch provisioning, see Add a Device Using the Serial Number (Zero-Touch Provisioning)—Basic Configuration, on page 41.

The **DDNS** page also supports setting DHCP server settings relating to DDNS.

Note DDNS is not supported on the BVI or bridge group member interfaces.

Before you begin

- Configure a DNS server group on **Objects** > **Object Management** > **DNS Server Group**, and then enable the group for the interface on **Devices** > **Platform Settings**, create or edit a Threat Defense policy and click **DNS**. See DNS, on page 949.
- Configure the device hostname. You can configure the hostname when you perform the Firewall Threat Defense initial setup, or by using the **configure network hostname** command. If you do not specify the hostname per interface, then the device hostname is used.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **DHCP** > **DDNS**.

Step 3 Standard DDNS method: Configure a DDNS update method to enable DNS requests from the Firewall Threat Defense.

You do not need to configure a DDNS update method if the DHCP server will perform all requests.

a) On **DDNS Update Methods**, click **Add**.
b) Set the **Method Name**.
c) Click **DDNS**.
d) (Optional) Configure the **Update Interval** between DNS requests. By default when all values are set to 0, update requests are sent whenever the IP address or hostname changes. To send requests regularly, set the **Days** (0-364), **Hours**, **Minutes**, and **Seconds**.
e) Set the **Update Records** you want the Firewall Threat Defense to update.

This setting only affects the records you want to update directly from the Firewall Threat Defense; to determine the records you want the DHCP server to update, configure the DHCP client settings per interface or globally. See, Step 5, on page 922.

- **Not Defined**—Disables DNS updates from the Firewall Threat Defense.

- **Both A and PTR Records**—Sets the Firewall Threat Defense to update both A and PTR RRs. Use this option for static or PPPoE IP addressing.
- **A Records**—Sets the Firewall Threat Defense to update the A RR only. Use this option if you want the DHCP server to update the PTR RR.

f) Click **OK**.
g) Assign this method to the interface in Step 5, on page 922.

Step 4 Web method: Configure a DDNS update method to enable HTTP update requests from the Firewall Threat Defense.

a) On **DDNS Update Methods**, click **Add**.
b) Set the **Method Name**.
c) Click **Web**.
d) Set the **Web Update Type** to update IPv4, IPv6, or both types of addresses.
e) Set the **Web URL**. Specify the update URL. Check with your DNS provider for the URL required.

Use the following syntax:

https://*username***:***password***@***provider-domain***/***path***?hostname=<h>&myip=<a>**

Example:

https://jcrichton:pa$$w0rd17@domains.example.com/nic/update?hostname=<h>&myip=<a>

f) (Optional) Configure the **Update Interval** between DNS requests. By default when all values are set to 0, update requests are sent whenever the IP address or hostname changes. To send requests regularly, set the **Days** (0-364), **Hours**, **Minutes**, and **Seconds**.
g) Click **OK**.
h) Assign this method to the interface in Step 5, on page 922.
i) The web type method for DDNS also requires you to identify the DDNS server root CA to validate the DDNS server certificate for the HTTPS connection. See, Step 9, on page 924.

Step 5 Configure interface settings for DDNS, including setting the update method, DHCP client settings, and the hostname for this interface.

a) On **DDNS Interface Settings**, click **Add**.
b) Choose the **Interface** from the drop-down list.
c) Choose the **Method Name** that you created on the **DDNS Update Methods** page.

(Standard DDNS method) You do not need to assign a method if you want the DHCP server to perform all updates.

d) Set the **Host Name** for this interface.

If you do not set the hostname, the device hostname is used. If you do not specify an FQDN, then the default domain from the DNS server group is appended (for static or PPPoE IP addressing) or the domain name from the DHCP server is appended (for DHCP IP addressing).

e) Standard DDNS method: Configure the **DHCP Client requests DHCP server to update requests** to determine which records you want the DHCP server to update.

The Firewall Threat Defense sends DHCP client requests to the DHCP server. Note that the DHCP server must also be configured to support DDNS. The server can be configured to honor the client requests, or it can override the client (in which case, it will reply to the client so the client does not also try to perform updates that the server is performing).

For static or PPPoE IP addressing, these settings are ignored.

Note
You can also set these values globally for all interfaces on the **DDNS** page. The per-interface settings take precedence over the global settings.

- **Not Selected**—Disables DDNS requests to the DHCP server. Even if the client does not request DDNS updates, the DHCP server can be configured to send updates anyway.
- **No Update**—Requests the DHCP server not to perform updates. This setting works in conjunction with a DDNS update method with **Both A and PTR Records** enabled.
- **Only PTR**—Requests that the DHCP server perform the PTR RR update. This setting works in conjunction with a DDNS update method with **A Records** enabled.
- **Both A and PTR Records**—Requests that the DHCP server perform both A and PTR RR updates. This setting does not require a DDNS update method to be associated with the interface.

f) Click **OK**.

Note
The **Dynamic DNS Update** settings relate to DHCP server settings when you enable a DHCP server on the Firewall Threat Defense. See, Step 6, on page 923 for more information.

Step 6 If you enable the DHCP server on an Firewall Threat Defense, you can configure DHCP server settings for DDNS.

To enable the DHCP server, see Configure the DHCPv4 Server, on page 913). You can configure the server behavior when DHCP clients use the standard DDNS update method. If the server performs any updates, then if the client lease expires (and is not renewed), the server will request that the DNS server remove the RRs for which it was responsible.

a) You can configure server settings globally or per interface. For global settings, see the main **DDNS** page. For per-interface settings, see the **DDNS Interface Settings** page. Interface settings take precedence over global settings.
b) Configure which DNS RRs you want the DHCP server to update under **Dynamic DNS Update**.

- **Not Selected**—DDNS updates are disabled, even if the client requests them.
- **Only PTR**—Enables DDNS updates. If you enable the **Override DHCP Client Requests** setting, then the server will only update the PTR RR. Otherwise, the server will update RRs that the client requests. If the client does not send an update request with the FQDN option, the server will request an update for both A and PTR RRs using the hostname discovered in DHCP option 12.
- **Both A and PTR Records**—Enables DDNS updates. If you enable the **Override DHCP Client Requests** setting, then the server will update both the A and PTR RRs. Otherwise, the server will update RRs that the client requests. If the client does not send an update request with the FQDN option, the server will request an update for both A and PTR RRs using the hostname discovered in DHCP option 12.

c) To override the update actions requested by the DHCP client, check **Override DHCP Client Requests**.

The server will reply to the client that the request was overridden, so the client does not also try to perform updates that the server is performing.

Step 7 (Optional) Configure general DHCP client settings. These settings are not related to DDNS, but are related to how the DHCP client behaves.

a) On the **DDNS** page, check **Enable DHCP Client Broadcast** to request that the DHCP server broadcast the DHCP reply (DHCP option 1).

b) To force a MAC address to be stored inside a DHCP request packet for option 61 instead of the default internally generated string, on **DDNS** > **DHCP Client ID Interface**, choose the interface from the **Available Interfaces** list, and then click **Add** to move it to the **Selected Interfaces** list.

Some ISPs expect option 61 to be the interface MAC address. If the MAC address is not included in the DHCP request packet, then an IP address will not be assigned. This setting does not directly relate to DDNS, but is a general DHCP client setting.

Step 8 Click **Save** on the Device page to save your changes.

Step 9 The Web method for DDNS also requires you to identify the DDNS server root CA to validate the DDNS server certificate for the HTTPS connection.

The following example shows how to add a DDNS server's CA as a trustpoint.

a) Obtain the DDNS server CA certificate. This procedure shows a manual import using PEM format, but you can also use PKCS12.

b) In Firewall Management Center, choose **Devices** > **Certificates**, and click **Add**.

c) Select a **Device**, and click **Add** (+).

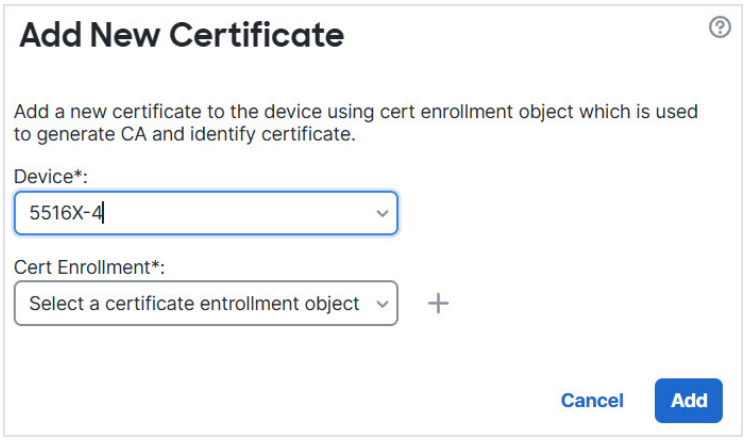

The **Add Cert Enrollment** dialog box appears.

d) Enter the following fields, and click **Save**:

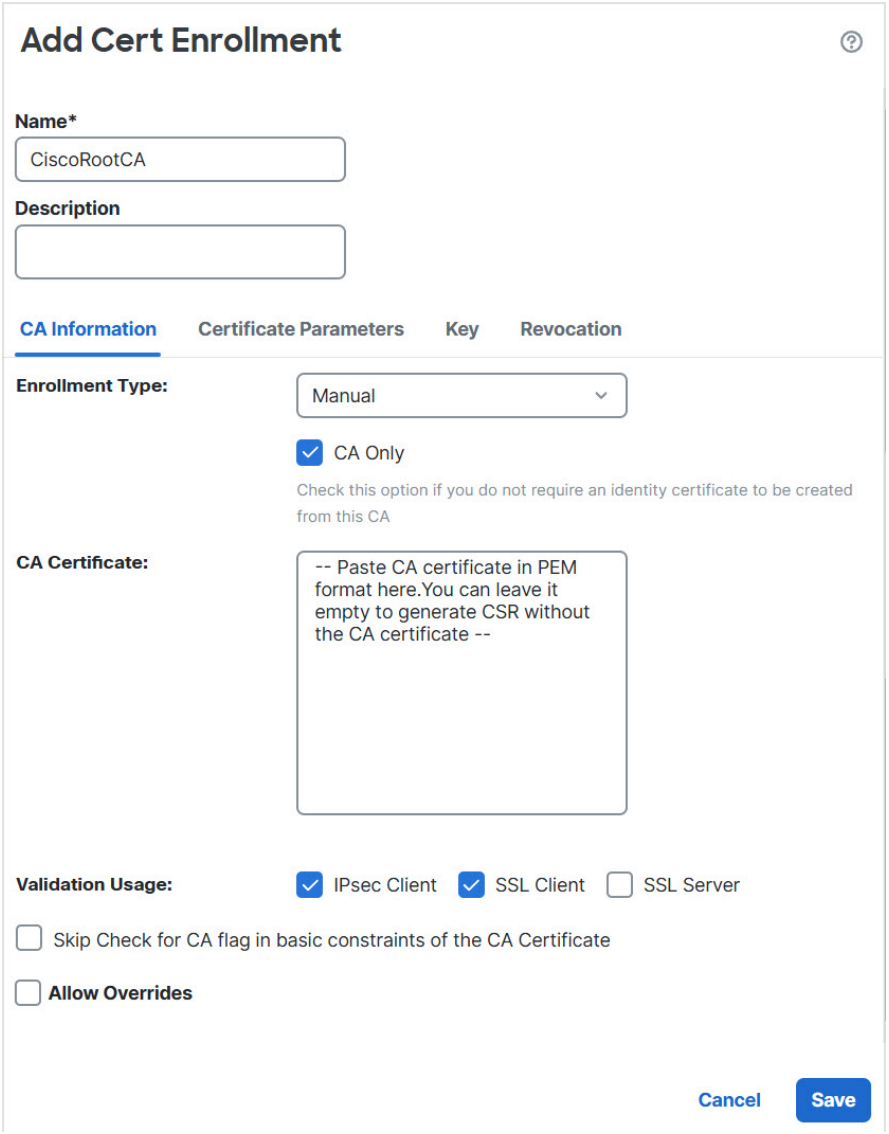

- Enter a **Name**.
- Choose **Enrollment Type** > **Manual**.
- Click **CA Only**.
- Paste in the CA text from step 9.a, on page 924.

e) Click **Save**.

History for DHCP and DDNS

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Configure DHCP relay trusted interfaces from the Firewall Management Center web interface.	7.2.6/7.4.1	Any	**Upgrade impact. Redo any related FlexConfigs after upgrade.** You can now use the Firewall Management Center web interface to configure interfaces as trusted interfaces to preserve DHCP Option 82. If you do this, these settings override any existing FlexConfigs, although you should remove them. DHCP Option 82 is used by downstream switches and routers for DHCP snooping and IP Source Guard. Normally, if the Firewall Threat Defense DHCP relay agent receives a DHCP packet with Option 82 already set, but the giaddr field (which specifies the DHCP relay agent address that is set by the relay agent before it forwards the packet to the server) is set to 0, then Firewall Threat Defense will drop that packet by default. You can preserve Option 82 and forward the packet by identifying an interface as a trusted interface. New/modified screens: **Devices** > **Device Management** > **Add/Edit Device** > **DHCP** > **DHCP Relay** Version restrictions: Not supported with Firewall Management Center Version 7.3.x or 7.4.0. If you upgrade to an unsupported version, redo your FlexConfigs.
DHCPv6 Stateless Server	7.3.0	7.3.0	The Firewall Threat Defense now supports a light DHCPv6 stateless server when using the DHCPv6 Prefix Delegation client. The Firewall Threat Defense provides other information such as the domain name to SLAAC clients when they send Information Request (IR) packets to the Firewall Threat Defense. The Firewall Threat Defense only accepts IR packets and does not assign addresses to the clients. New/modified screens: • **Devices** > **Device Management** > **Interfaces** > **Add/Edit Interfaces** > **IPv6** > **DHCP** • **Objects** > **Object Management** > **DHCP IPv6 Pool** New/modified commands: **show ipv6 dhcp**

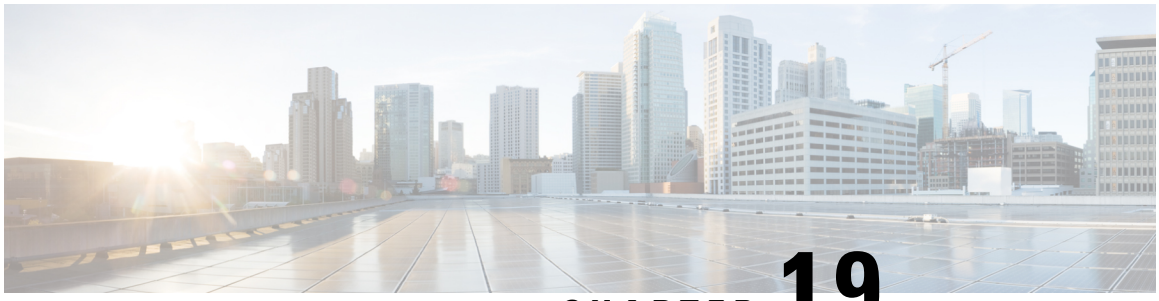

CHAPTER 19

SNMP for the Firepower 1000

This chapter describes how to configure SNMP for the Firepower 1000.

- About SNMP for the Firepower 1000, on page 927
- Enabling SNMP and Configuring SNMP Properties for Firepower 1000, on page 927
- Creating an SNMP Trap for Firepower 1000, on page 929
- Creating an SNMP User for Firepower 1000, on page 930

About SNMP for the Firepower 1000

The Simple Network Management Protocol (SNMP) is an application-layer protocol that provides a message format for communication between SNMP managers and agents. SNMP provides a standardized framework and a common language used for the monitoring and management of devices in a network.

The SNMP framework consists of three parts:

- An SNMP manager—The system used to control and monitor the activities of network devices using SNMP.
- An SNMP agent—The software component within the Firepower 1000 chassis that maintains the data for the Firepower chassis and reports the data, as needed, to the SNMP manager. The Firepower chassis includes the agent and a collection of MIBs. To enable the SNMP agent and create the relationship between the manager and agent, enable and configure SNMP in the Firewall Management Center.
- A managed information base (MIB)—The collection of managed objects on the SNMP agent.

The Firepower 1000 chassis supports SNMPv1, SNMPv2c and SNMPv3. Both SNMPv1 and SNMPv2c use a community-based form of security; SNMPv3 use username and password for security.

Enabling SNMP and Configuring SNMP Properties for Firepower 1000

Note This procedure only applies to the Firepower 1000.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **SNMP**.

Step 3 Complete the following fields:

Name	Description
Admin State check box	Whether SNMP is enabled or disabled. Enable this service only if your system includes integration with an SNMP server.
Port field	The port on which the Firepower chassis communicates with the SNMP host. You cannot change the default port.
Community field	The default SNMPv1 or v2 community name or SNMP v3 username and password that the Firepower chassis includes on any trap messages it sends to the SNMP host. Enter a valid community string for SNMPv1 and SNMPv2: • Alphanumeric string between 1 and 32 characters and special characters ! (exclamation), - (hyphen), ~ (tilde), && (double ampersand), [ ] (square brackets), ^ (carat), ' (single quote), " (double quotes), and < > (angle brackets). • Do not use @ (at sign), \ (backslash), ? (question mark) or an empty space. • The string can also be in ASCII characters ranging 0x21 to 0x7E inclusive, excluding HTML interjection vectors, namely single quote ('), double quotes ("), and angle brackets (< >). Enter a valid username and password for SNMPv3: • Username can be alphanumeric string and can include @ (at sign), \ (backslash), . (period), _ (underscore), and - (hyphen). • The password restrictions are same as the community string restrictions. Note that if the **Community** field is already set, the text to the right of the empty field reads **Set: Yes**. If the **Community** field is not yet populated with a value, the text to the right of the empty field reads **Set: No**.
System Admin Name field	The contact person responsible for the SNMP implementation. Enter a string of up to 255 characters, such as an email address or a name and telephone number.
Location field	The location of the host on which the SNMP agent (server) runs. Enter an alphanumeric string up to 510 characters.

Step 4 Click **Save**.

What to do next

Create SNMP traps and users.

Creating an SNMP Trap for Firepower 1000

Note This procedure only applies to the Firepower 1000.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **SNMP**.

Step 3 In the **SNMP Traps Configuration** area, click **Add**.

Step 4 In the **SNMP Trap Configuration** dialog box, complete the following fields:

Name	Description
Host Name field	The hostname or IP address of the SNMP host to which the Firepower chassis should send the trap.
Community field	The SNMP v1 or v2 community name or the SNMP v3 username the Firepower chassis includes when it sends the trap to the SNMP host. This must be the same as the community or username that is configured for the SNMP service. Enter an alphanumeric string between 1 and 32 characters. Do not use @ (at sign), \ (backslash), " (double quote), ? (question mark) or an empty space.
Port field	The port on which the Firepower chassis communicates with the SNMP host for the trap. Enter an integer between 1 and 65535.
Version field	The SNMP version and model used for the trap. This can be one of the following: • **V1** • **V2** • **V3**

Name	Description
Type field	If you select **V2** or **V3** for the version, the type of trap to send. This can be one of the following: • **Traps** • **Informs**
Privilege field	If you select **V3** for the version, the privilege associated with the trap. This can be one of the following: • **Auth**—Authentication but no encryption • **Noauth**—No authentication or encryption • **Priv**—Authentication and encryption

Step 5 Click **OK** to close the **SNMP Trap Configuration** dialog box.

Step 6 Click **Save**.

Creating an SNMP User for Firepower 1000

Note This procedure only applies to the Firepower 1000.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click **SNMP**.

Step 3 In the **SNMP Users Configuration** area, click **Add**.

Step 4 In the **SNMP User Configuration** dialog box, complete the following fields:

Name	Description
Username field	The username assigned to the SNMP user. Enter up to 32 letters or numbers. The name must begin with a letter and you can also specify _ (underscore), . (period), @ (at sign), and - (hyphen).
Auth Algorithm Type field	The authorization type: **SHA**.

Name	Description
Use AES-128 checkbox	If checked, this user uses AES-128 encryption. **Note** SNMPv3 does not support DES. If you leave the AES-128 box unchecked, no privacy encryption will be done and any configured privacy password will have no effect.
Authentication Password field	The password for the user.
Confirm field	The password again for confirmation purposes.
Encryption Password field	The privacy password for the user.
Confirm field	The privacy password again for confirmation purposes.

Step 5 Click **OK** to close the **SNMP User Configuration** dialog box.

Step 6 Click **Save**.

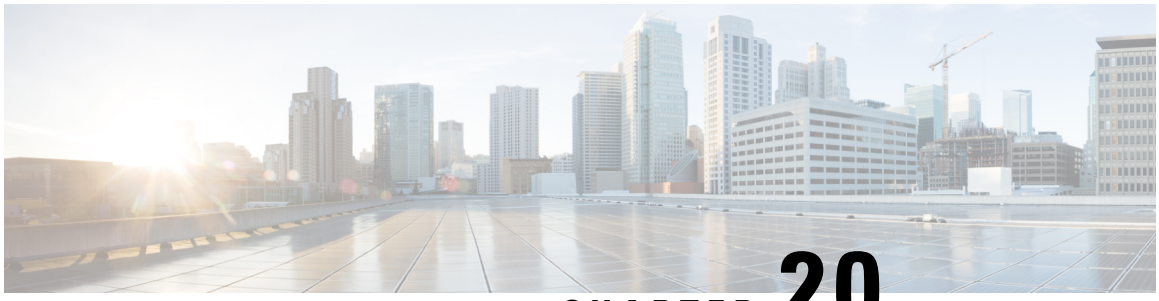

CHAPTER 20

Quality of Service

The following topics describe how to use the Quality of Service (QoS) feature to police network traffic using Firewall Threat Defense devices:

- Introduction to QoS, on page 933
- About QoS Policies, on page 933
- Requirements and Prerequisites for QoS, on page 934
- Rate Limiting with QoS Policies, on page 934
- History for QoS, on page 943

Introduction to QoS

Quality of Service, or QoS, rate limits (polices) network traffic that is allowed or trusted by access control. The system does not rate limit traffic that was fastpathed.

Though QoS is supported only on the routed interfaces of Firewall Threat Defense devices, it is not supported on site-to-site VPN and VTI interfaces.

Logging Rate-Limited Connections

There are no logging configurations for QoS. A connection can be rate limited without being logged, and you cannot log a connection simply because it was rate limited. To view QoS information in connection events, you must independently log the ends of the appropriate connections to the Firewall Management Center database. See *Other Connections You Can Log* in the Cisco Secure Firewall Management Center Administration Guide for more information.

Connection events for rate-limited connections contain information on how much traffic was dropped, and which QoS configurations limited the traffic. You can view this information in event views (workflows), dashboards, and reports.

About QoS Policies

QoS policies deployed to managed devices govern rate limiting. Each QoS policy can target multiple devices; each device can have one deployed QoS policy at a time.

The system matches traffic to QoS rules in the order you specify. The system rate limits traffic according to the first rule where all rule conditions match the traffic. Traffic that does not match any of the rules is not rate limited.

Note The total number of rules including QoS rules on the device cannot exceed 255. When this threshold is reached, a deployment warning message is displayed. You need to reduce the number of rules for a successful deployment.

You must constrain QoS rules by source or destination (routed) interfaces. The system enforces rate limiting *independently* on *each* of those interfaces; you cannot specify an aggregate rate limit for a set of interfaces.

QoS rules can also rate limit traffic by other network characteristics, as well as contextual information such as application, URL, user identity, and custom Security Group Tags (SGTs).

You can rate limit download and upload traffic independently. The system determines download and upload directions based on the connection initiator.

Note QoS is not subordinate to a main access control configuration; you configure QoS independently. However, the access control and QoS policies deployed to the same device share identity configurations; see Associating Other Policies with Access Control, on page 1762.

Requirements and Prerequisites for QoS

Model Support

Threat Defense

Supported Domains

Any

User Roles

Admin

Access Admin

Network Admin

Rate Limiting with QoS Policies

To perform policy-based rate limiting, configure and deploy QoS policies to managed devices. Each QoS policy can target multiple devices; each device can have one deployed QoS policy at a time.

Only one person should edit a policy at a time, using a single browser window. If multiple users save the same policy, the last saved changes are retained. For your convenience, the system displays information on who (if anyone) is currently editing each policy. To protect the privacy of your session, a warning appears after 30 minutes of inactivity on the policy editor. After 60 minutes, the system discards your changes.

Procedure

Step 1 Choose **Devices** > **QoS**.

Step 2 Click **New Policy** to create a new QoS policy and, optionally, assign target devices; see Creating a QoS Policy, on page 935.

You can also **Copy** ( ) or **Edit** ( ) an existing policy.

Step 3 Configure QoS rules; see Configuring QoS Rules, on page 936 and QoS Rule Conditions, on page 938.

The Rules in the QoS policy editor lists each rule in evaluation order, and displays a summary of the rule conditions and rate limiting configurations. A right-click menu provides rule management options, including moving, enabling, and disabling.

Helpful in larger deployments, you can **Filter by Device** to display only the rules that affect a specific device or group of devices. You can also search for and within rules; the system matches text you enter in the **Search Rules** field to rule names and condition values, including objects and object groups.

Note
Properly creating and ordering rules is a complex task, but one that is essential to building an effective deployment. If you do not plan carefully, rules can preempt other rules, require additional licenses, or contain invalid configurations. Icons represent comments, warnings, and errors. If issues exist, click **Show Warnings** to display a list. For more information, see Best Practices for Access Control Rules, on page 1738.

Step 4 Click **Policy Assignments** to identify the managed devices targeted by the policy; see Setting Target Devices for a QoS Policy, on page 936.

If you identified target devices during policy creation, verify your choices.

Step 5 Save the QoS policy.

Step 6 Because this feature must allow some packets to pass, you must configure your system to examine those packets..

Step 7 Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Creating a QoS Policy

A new QoS policy with no rules performs no rate limiting.

Procedure

Step 1 Choose **Devices** > **QoS**.

Step 2 Click **New Policy**.

Step 3 Enter a **Name** and, optionally, a **Description**.

Step 4 (Optional) Choose the **Available Devices** where you want to deploy the policy, then click **Add to Policy**, or drag and drop to the **Selected Devices**. To narrow the devices that appear, type a search string in the **Search** field.

You must assign devices before you deploy the policy.

Step 5 Click **Save**.

What to do next

- Configure and deploy the QoS policy; see Rate Limiting with QoS Policies, on page 934.

Setting Target Devices for a QoS Policy

Each QoS policy can target multiple devices; each device can have one deployed QoS policy at a time.

Procedure

Step 1 In the QoS policy editor, click **Policy Assignments**.

Step 2 Build your target list:

- Add—Choose one or more **Available Devices**, then click **Add to Policy** or drag and drop into the list of **Selected Devices**.
- Delete—Click **Delete** ( ) next to a single device, or choose multiple devices, right-click, then choose **Delete Selected**.
- Search—Enter a search string in the search field. Click **Clear** ( ) to clear the search.

Step 3 Click **OK** to save policy assignments.

Step 4 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Configuring QoS Rules

When you create or edit a rule, use the upper portion of the rule editor to configure general rule properties. Use the lower portion of the rule editor to configure rule conditions and comments.

Procedure

Step 1 On Rules of the QoS policy editor:

- Add Rule—Click **Add Rule**.
- Edit Rule—Click **Edit** ( ).

Step 2 Enter a **Name**.

Step 3 Configure rule components:

- Enabled—Specify whether the rule is **Enabled**.
- Apply QoS On—Choose the interfaces you want to rate limit, either **Interfaces in Destination Interface Objects** or **Interfaces in Source Interface Objects**. Your choice must correspond with a populated interface constraint (not **any**).
- Traffic Limit Per Interface—Enter a **Download Limit** and an **Upload Limit** in Mbits/sec. The default value of **Unlimited** prevent matching traffic from being rate limited in that direction.
- Conditions—Click the corresponding condition you want to add. You must configure a source or destination interface condition, corresponding to your choice for **Apply QoS On**.
- Comments—Click **Comments**. To add a comment click **New Comment**, enter a comment, and click **OK**. You can edit or delete this comment until you save the rule.

For detailed information on rule components, see QoS Rule Components, on page 937.

Step 4 Save the rule.

Step 5 In the policy editor, set the rule position. Click and drag or use the right-click menu to cut and paste.

Rules are numbered starting at 1. The system matches traffic to rules in top-down order by ascending rule number. The first rule that traffic matches is the rule that handles that traffic. Proper rule order reduces the resources required to process network traffic and prevents rule preemption.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Related Topics

Best Practices for Access Control Rules, on page 1738

QoS Rule Components

State (Enabled/Disabled)

By default, rules are enabled. If you disable a rule, the system does not use it and stops generating warnings and errors for that rule.

Interfaces (Apply QoS On)

You cannot save a QoS rule that rate limits all traffic. For each QoS rule, you must apply QoS on either:

- Interfaces in Source Interface Objects—Rate limits traffic through the rule's source interfaces. If you choose this option, you must add at least one source interface constraint (cannot be **any**).
- Interfaces in Destination Interface Objects—Rate limits traffic through the rule's destination interfaces. If you choose this option, you must add at least one destination interface constraint (cannot be **any**).

Traffic Limit Per Interface

A QoS rule enforces rate limiting *independently* on *each* of the interfaces you specify with the Apply QoS On option. You cannot specify an aggregate rate limit for a set of interfaces.

You can rate limit traffic by Mbits per second. The default value of **Unlimited** prevents matching traffic from being rate limited.

You can rate limit download and upload traffic independently. The system determines download and upload directions based on the connection initiator.

If you specify a limit greater than the maximum throughput of an interface, the system does not rate limit matching traffic. Maximum throughput may be affected by an interface's hardware configuration, which you specify in each device's properties (**Devices** > **Device Management**).

Conditions

Conditions specify the specific traffic the rule handles. You can configure each rule with multiple conditions. Traffic must match all conditions to match the rule. Each condition type has its own tab in the rule editor. For more information, see QoS Rule Conditions, on page 938.

Comments

Each time you save changes to a rule you can add comments. For example, you might summarize the overall configuration for the benefit of other users, or note when you change a rule and the reason for the change.

In the policy editor, the system displays how many comments a rule has. In the rule editor, use the Comments tab to view existing comments and add new ones.

QoS Rule Conditions

Conditions specify the specific traffic the rule handles. You can configure each rule with multiple conditions. Traffic must match all conditions to match the rule. Each condition type has its own tab in the rule editor. You can rate limit traffic using:

See one of the following sections for more information.

Related Topics

Interface Rule Conditions, on page 938
Network Rule Conditions, on page 939
User Rule Conditions, on page 939
Application Rule Conditions, on page 939
Port Rule Conditions, on page 940
URL Rule Conditions, on page 942
Custom SGT Rule Conditions, on page 942

Interface Rule Conditions

Interface rule conditions control traffic by its source and destination interfaces.

Depending on the rule type and the devices in your deployment, you can use predefined *interface objects* called *security zones* or *interface groups* to build interface conditions. Interface objects segment your network to help you manage and classify traffic flow by grouping interfaces across multiple devices; see Interface, on page 1398. QoS rules can be applied only on routed interfaces.

Tip Constraining rules by interface is one of the best ways to improve system performance. If a rule excludes all of a device's interfaces, that rule does not affect that device's performance.

Network Rule Conditions

Networks control or decrypt traffic by its source and destination IP address, using inner headers. Tunnel rules, which use outer headers, have tunnel endpoint conditions instead of network conditions.

You can use predefined objects to build network conditions, or manually specify individual IP addresses or address blocks.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Note You *cannot* use FDQN network objects in identity rules.

User Rule Conditions

Matches traffic based the user who initiates the connection, or the group to which the user belongs. For example, you could configure a Block rule to prohibit anyone in the Finance group from accessing a network resource.

You can configure user rule conditions for users in Microsoft Active Directory realms only.

In addition to configuring users and groups for configured realms, you can set policies for the following Special Identities users:

- Failed Authentication: User that failed authentication with the captive portal.
- Guest: Users configured as guest users in the captive portal.
- No Authentication Required: Users that match an identity **No Authentication Required** rule action.
- Unknown: Users that cannot be identified; for example, users that are not downloaded by a configured realm.

For access control rules only, you must first associate an identity policy with the access control policy as discussed in Associating Other Policies with Access Control, on page 1762.

Application Rule Conditions

When the system analyzes IP traffic, it can identify and classify the commonly used applications on your network. This discovery-based *application awareness* is the basis for *application control*—the ability to control application traffic.

System-provided *application filters* help you perform application control by organizing applications according to basic characteristics: type, risk, business relevance, category, and tags. You can create reuseable user-defined filters based on combinations of the system-provided filters, or on custom combinations of applications.

At least one detector must be enabled for each application rule condition in the policy. If no detector is enabled for an application, the system automatically enables all system-provided detectors for the application; if none exist, the system enables the most recently modified user-defined detector for the application. For more information about application detectors, see Application Detector Fundamentals, on page 2624.

You can use both application filters and individually specified applications to ensure complete coverage. However, understand the following note before you order your access control rules.

Benefits of Application Filters

Application filters help you quickly configure application control. For example, you can easily use system-provided filters to create an access control rule that identifies and blocks all high risk, low business relevance applications. If a user attempts to use one of those applications, the system blocks the session.

Using application filters simplifies policy creation and administration. It assures you that the system controls application traffic as expected. Because Cisco frequently updates and adds application detectors via system and vulnerability database (VDB) updates, you can ensure that the system uses up-to-date detectors to monitor application traffic. You can also create your own detectors and assign characteristics to the applications they detect, automatically adding them to existing filters.

Application Characteristics

The system characterizes each application that it detects using the criteria described in the following table. Use these characteristics as application filters.

Table 54: Application Characteristics

Characteristic	Description	Example
Type	Application protocols represent communications between hosts. Clients represent software running on a host. Web applications represent the content or requested URL for HTTP traffic.	HTTP and SSH are application protocols. Web browsers and email clients are clients. MPEG video and Facebook are web applications.
Risk	The likelihood that the application is being used for purposes that might be against your organization's security policy.	Peer-to-peer applications tend to have a very high risk.
Business Relevance	The likelihood that the application is being used within the context of your organization's business operations, as opposed to recreationally.	Gaming applications tend to have a very low business relevance.
Category	A general classification for the application that describes its most essential function. Each application belongs to at least one category.	Facebook is in the social networking category.
Tag	Additional information about the application. Applications can have any number of tags, including none.	Video streaming web applications often are tagged `high bandwidth` and `displays ads`.

Related Topics

Best Practices for Configuring Application Control, on page 1735

Port Rule Conditions

Port conditions allow you to control traffic by its source and destination ports.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Best Practices for Port-Based Rules

Specifying ports is the traditional way to target applications. However, applications can be configured to use unique ports to bypass access control blocks. Thus, whenever possible, use application filtering criteria rather than port criteria to target traffic.

Application filtering is also recommended for applications, like Firewall Threat Defense, that open separate channels dynamically for control vs. data flow. Using port-based access control rules can prevent these kinds of applications from performing correctly, and could result in blocking desirable connections.

Using Source and Destination Port Constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as source port conditions in a single access control rule.

Port, Protocol, and ICMP Code Rule Conditions

Port conditions match traffic based on the source and destination ports. Depending on the rule type, "port" can represent any of the following:

- TCP and UDP—You can control TCP and UDP traffic based on the port. The system represents this configuration using the protocol number in parentheses, plus an optional associated port or port range. For example: TCP(6)/22.
- ICMP—You can control ICMP and ICMPv6 (IPv6-ICMP) traffic based on its internet layer protocol plus an optional type and code. For example: ICMP(1):3:3.
- Protocol—You can control traffic using other protocols that do not use ports.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Best Practices for Port-Based Rules

Specifying ports is the traditional way to target applications. However, applications can be configured to use unique ports to bypass access control blocks. Thus, whenever possible, use application filtering criteria rather than port criteria to target traffic. Note that application filtering is not available in prefilter rules.

Application filtering is also recommended for applications, like FTP, that open separate channels dynamically for control vs. data flow. Using port-based access control rules can prevent these kinds of applications from performing correctly, and could result in blocking desirable connections.

Using Source and Destination Port Constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as destination port conditions in a single access control rule.

Matching Non-TCP Traffic with Port Conditions

You can match non-port-based protocols. By default, if you do not specify a port condition, you are matching IP traffic. Although you can configure port conditions to match non-TCP traffic, there are some restrictions:

- Access control rules—For Classic devices, you can match GRE-encapsulated traffic with an access control rule by using the GRE (47) protocol as a destination port condition. To a GRE-constrained rule, you can add only network-based conditions: zone, IP address, port, and VLAN tag. Also, the system uses outer headers to match **all** traffic in access control policies with GRE-constrained rules. For Firewall Threat Defense devices, use tunnel rules in the prefilter policy to control GRE-encapsulated traffic.
- Decryption rules—These rules support TCP port conditions only.
- IMCP echo—A destination ICMP port with the type set to 0 or a destination ICMPv6 port with the type set to 129 only matches unsolicited echo replies. ICMP echo replies sent in response to ICMP echo requests are ignored. For a rule to match on any ICMP echo, use ICMP type 8 or ICMPv6 type 128.

URL Rule Conditions

Use URL conditions to control the websites that users on your network can access.

For complete information, see About URL Filtering with Category and Reputation, on page 1858.

Custom SGT Rule Conditions

If you do not configure ISE/ISE-PIC as an identity source, you can control traffic using Security Group Tags (SGTs) that were **not** assigned by ISE. SGTs specify the privileges of traffic sources within a trusted network.

Custom SGT rule conditions use manually created SGT objects to filter traffic, rather than ISE SGTs obtained from the system's connection to an ISE server. These manually created SGT objects correspond to the SGT attributes on the traffic you want to control. Controlling traffic using custom SGTs is not considered user control.

ISE SGT vs Custom SGT Rule Conditions

Some rules allow you to control traffic based on assigned SGT. Depending on the rule type and your identity source configuration, you can use either ISE-assigned SGTs or custom SGTs to match traffic with assigned SGT attributes.

Note If you use ISE SGTs to match traffic, even if a packet does not have an assigned SGT attribute, the packet still matches an ISE SGT rule if the SGT associated with the packet's source IP address is known in ISE.

Condition Type	Requires	SGTs Listed in Rule Editor
ISE SGT	ISE identity source	SGTs obtained by querying the ISE server, with automatically updated metadata

Condition Type	Requires	SGTs Listed in Rule Editor
Custom SGT	No ISE/ISE-PIC identity source	Static SGT objects you create

Autotransition from Custom SGTs to ISE SGTs

If you create rules that match custom SGTs, then configure ISE/ISE-PIC as an identity source, the system:

- Disables **Security Group Tag** options in the object manager. Although the system retains existing SGT objects, you cannot modify them or add new ones.
- Retains existing rules with custom SGT conditions. However, these rules do not match traffic. You also cannot add additional custom SGT criteria to existing rules, or create new rules with custom SGT conditions.

If you configure ISE, Cisco recommends that you delete or disable existing rules with custom SGT conditions. Instead, use ISE attribute conditions to match traffic with SGT attributes.

History for QoS

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Deprecated: priority-queue with FlexConfig.	7.2.5	7.2.5	FlexConfig was used to configure priority-queue in threat defense. This command was removed.
Ability to specify handling of URLs having unknown reputation.	6.7.0	Any	For details, see History for URL Filtering, on page 1882. New/modified screens: QoS rule editor
Rate limit increased.	6.2.1	Any	Raised the maximum rate limit from 1,000 Mbps to 100,000 Mbps. New/modified screens: QoS rule editor
Custom SGT and original client network filtering.	6.2.1	Any	Rate limit traffic using custom Security Group Tags (SGTs) and original client network information (XFF, True-Client-IP, or custom-defined HTTP headers). New/modified screens: QoS rule editor
QoS (rate limiting) introduced.	6.1.0	Any	FTD can rate limit (police) network traffic that is allowed or trusted by access control. New/modified screens: **Devices** > **QoS**

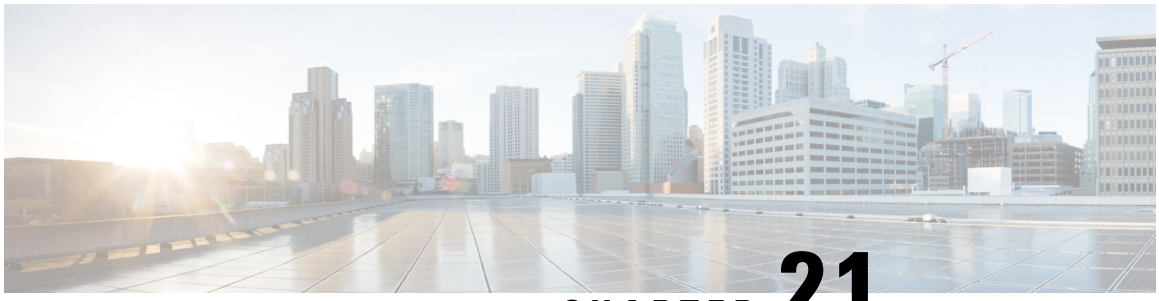

CHAPTER 21

Platform Settings

Platform settings for Firewall Threat Defense devices configure a range of unrelated features whose values you might want to share among several devices. Even if you want different settings per device, you must create a shared policy and apply it to the desired device.

Note In 7.4 and later, the Management and Diagnostic interfaces are merged. If Platform Settings for syslog servers or SNMP hosts specify the Diagnostic interface by name, then you must use separate Platform Settings policies for merged and non-merged devices (7.3 and earlier, and some upgraded 7.4 devices).

- Introduction to Platform Settings, on page 946
- Requirements and Prerequisites for Platform Settings Policies, on page 946
- Manage Platform Settings Policies, on page 946
- Chassis Platform Settings, on page 947
- ARP Inspection, on page 947
- Banner, on page 948
- DNS, on page 949
- External Authentication, on page 952
- Enable Virtual-Router-Aware Interface for External Authentication of Platform, on page 957
- Fragment Settings, on page 958
- **HTTP Access**, on page 959
- **ICMP Access**, on page 960
- NetFlow, on page 962
- **SSH Access**, on page 964
- SMTP Server, on page 966
- SNMP, on page 966
- SSL , on page 980
- Syslog, on page 984
- Timeouts, on page 1000
- Time Synchronization, on page 1002
- Time Zone, on page 1003
- UCAPL/CC Compliance, on page 1004
- Performance Profile, on page 1005
- History for Platform Settings, on page 1006

Introduction to Platform Settings

A platform settings policy is a shared set of features or parameters that define the aspects of a managed device that are likely to be similar to other managed devices in your deployment, such as time settings and external authentication.

A shared policy makes it possible to configure multiple managed devices at once, which provides consistency in your deployment and streamlines your management efforts. Any changes to a platform settings policy affects all the managed devices where you applied the policy. Even if you want different settings per device, you must create a shared policy and apply it to the desired device.

For example, your organization's security policies may require that your appliances have a "No Unauthorized Use" message when a user logs in. With platform settings, you can set the login banner once in a platform settings policy.

You can also benefit from having multiple platform settings policies on a single Firewall Management Center. For example, if you have different mail relay hosts that you use under different circumstances or if you want to test different access lists, you can create several platform settings policies and switch between them, rather than editing a single policy.

Requirements and Prerequisites for Platform Settings Policies

Supported Domains

Any

User Roles

Admin

Access Admin

Network Admin

Manage Platform Settings Policies

Use the **Platform Settings** page (**Devices** > **Platform Settings**) to manage platform settings policies. This page indicates the type of device for each policy. The **Status** column shows the device targets for the policy.

Procedure

Step 1 Choose **Devices** > **Platform Settings**.

Step 2 For an existing policy, you can **Copy** (), **Edit** (), or **Delete** () the policy.

Caution
You should not delete a policy that is the last-deployed policy on any of its target devices, even if it is out of date. Before you delete the policy completely, it is good practice to deploy a different policy to those targets.

Step 3 To create a new policy, click **New Policy**.

a) Choose a device type from the drop-down list:

- **Threat Defense Settings** to create a shared policy for managed Firewall Threat Defense devices.
- **Chassis Platform Settings** to create a shared policy for managed Firewall Threat Defense chassis in multi-instance mode.

b) Enter a **Name** for the new policy and optionally, a **Description**.

c) Optionally, choose the **Available Devices** or **Available Chassis** where you want to apply the policy and click **Add** (or drag and drop) to add the selected devices. You can enter a search string in the **Search** field to narrow the list of devices.

d) Click **Save**.

The system creates the policy and opens it for editing.

Step 4 To change the target devices for a policy, click **Edit** (✎) next to the platform settings policy that you want to edit.

a) Click **Policy Assignment**.

b) To assign a device, high-availability pair, or device group to the policy, select it in the **Available Devices** or **Available Chassis** list and click **Add**. You can also drag and drop.

c) To remove a device assignment, click **Delete** (🗑) next to a device, high-availability pair, or device group in the **Selected Devices** or **Selected Chassis** list.

d) Click **OK**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Chassis Platform Settings

Chassis Platform Settings apply to multiple-instance mode chassis. For details on these settings, see Configure Chassis Platform Settings, on page 395.

ARP Inspection

By default, all ARP packets are allowed between bridge group members. You can control the flow of ARP packets by enabling ARP inspection.

ARP inspection prevents malicious users from impersonating other hosts or routers (known as ARP spoofing). ARP spoofing can enable a “man-in-the-middle” attack. For example, a host sends an ARP request to the gateway router; the gateway router responds with the gateway router MAC address. The attacker, however, sends another ARP response to the host with the attacker MAC address instead of the router MAC address. The attacker can now intercept all the host traffic before forwarding it on to the router.

ARP inspection ensures that an attacker cannot send an ARP response with the attacker MAC address, so long as the correct MAC address and the associated IP address are in the static ARP table.

When you enable ARP inspection, the Firewall Threat Defense device compares the MAC address, IP address, and source interface in all ARP packets to static entries in the ARP table, and takes the following actions:

- If the IP address, MAC address, and source interface match an ARP entry, the packet is passed through.
- If there is a mismatch between the MAC address, the IP address, or the interface, then the Firewall Threat Defense device drops the packet.
- If the ARP packet does not match any entries in the static ARP table, then you can set the Firewall Threat Defense device to either forward the packet out all interfaces (flood), or to drop the packet.

Note The dedicated Management interface never floods packets even if this parameter is set to flood.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **ARP Inspection**.

Step 3 Add entries to the ARP inspection table.

a) Click **Add** to create a new entry, or click **Edit** if the entry already exists.

b) Select the desired options.

- **Inspect Enabled**—To perform ARP inspection on the selected interfaces and zones.
- **Flood Enabled**—Whether to flood ARP requests that do not match static ARP entries out all interfaces other than the originating interface or the dedicated management interface. This is the default behavior.

  If you do not elect to flood ARP requests, then only those requests that exactly match static ARP entries are allowed.
- **Security Zones**—Add the zones that contain the interfaces on which to perform the selected actions. The zones must be switched zones. For interfaces not in a zone, you can type the interface name into the field below the Selected Security Zone list and click **Add**. These rules will be applied to a device only if the device includes the selected interfaces or zones.

c) Click **OK**.

Step 4 Add static ARP entries according to Add a Static ARP Entry, on page 882.

Step 5 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Banner

You can configure messages to show users when they connect to the device command line interface (CLI).

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Banner**.

Step 3 Configure the banner.

Following are some tips and requirements for banners.

- Only ASCII characters are allowed. You can use line returns (press Enter), but you cannot use tabs.
- You can dynamically add the hostname or domain name of the device by including the variables **$(hostname)** or **$(domain)**.
- Although there is no absolute length restriction on banners, Telnet or SSH sessions will close if there is not enough system memory available to process the banner messages.
- From a security perspective, it is important that your banner discourage unauthorized access. Do not use the words "welcome" or "please," as they appear to invite intruders in. The following banner sets the correct tone for unauthorized access:

```
You have logged in to a secure device.
If you are not authorized to access this device,
log out immediately or risk criminal charges.
```

Step 4 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

DNS

The Domain Name System (DNS) servers are used to resolve hostnames to IP addresses. There are two DNS server settings that apply to different types of traffic: data and special management traffic. Data traffic includes any services that use FQDNs for which a DNS lookup is necessary, such as access control rules and remote access VPN. Special management traffic includes traffic originating on the Management interface such as configuration and database updates. This procedure only applies to *data* DNS servers. For *management* DNS settings, see the CLI **configure network dns servers** and **configure network dns searchdomains** commands.

To determine the correct interface for DNS server communications, the managed device uses a routing lookup, but which routing table is used depends on the interfaces for which you enable DNS. See the interface settings below for more information.

You can optionally configure multiple DNS server groups and use them to resolve different DNS domains. For example, you could have a catch-all default group that uses public DNS servers, for use with connections to the Internet. You could then configure a separate group to use internal DNS servers for internal traffic, for example, any connection to a machine in the example.com domain. Thus, connections to an FQDN using your organization’s domain name would be resolved using your internal DNS servers, whereas connections

to public servers use external DNS servers. These resolutions are used by any feature that uses data DNS resolution, such as NAT and access control rules.

You can configure trusted DNS services for DNS snooping using the Trusted DNS Servers tab. DNS snooping is used to map the application domains to IPs in order to detect the application on the first packet. Apart from configuring the trusted DNS servers, you can include the already configured servers in DNS group, DHCP pool, DHCP relay and DHCP client as trusted DNS servers.

Note For an application-based PBR, you must configure trusted DNS servers. You must also ensure that the DNS traffic passes through Firewall Threat Defense in a clear-text format (encrypted DNS is not supported) so that domains can be resolved to detect applications.

Before you begin

- Ensure you have created one or more DNS server groups. For more information, see Creating DNS Server Group Objects, on page 1384.
- Ensure you have created interface objects to connect to the DNS servers.
- Ensure that the managed device has appropriate static or dynamic routes to access the DNS servers.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit a Threat Defense policy.

Step 2 Click **DNS**.

Step 3 Click the **DNS Settings** tab.

Step 4 Check **Enable DNS name resolution by device**.

Step 5 Configure the DNS server groups.

a) Do any of the following in the DNS server group list:

- To add a group to the list, click **Add**. You cannot add another group once there are 30 filter domains configured within the existing list of server groups.
- To edit the settings for a group, click **Edit** (✎) next to the group.
- To remove a group, click **Delete** (🗑) next to the group. Removing a group does not delete the DNS server group object, it simply removes it from this list.

b) When adding or editing a group, configure the following settings, then click **OK**:

- **Select DNS Group**—Select an existing DNS server group object, or click **+** to create a new one.
- **Make as default**—Select this option to make this group the default group. Any DNS resolution request that does not match the filters for other groups will be resolved using the servers in this group.
- **Filter Domains**—For non-default groups only, a comma-separated list of domain names, such as example.com,example2.com. Do not include spaces.

The group will be used for DNS resolutions for these domains only. You can enter a maximum of 30 separate domains across all groups added to this DNS platform settings policy. Each name can be a maximum of 127 characters.

Note that these filter domains are not related to the default domain name for the group. The filter list can be different from the default domain.

Step 6 (Optional) Enter the **Expiry Entry Timer** and **Poll Timer** values in minutes.

These options apply to FQDNs that are specified in network objects only. These do not apply to FQDNs used in other features.

- **Expire Entry Timer** specifies the minimum time-to-live (TTL) for the DNS entry, in minutes. If the expiration timer is longer than the entry's TTL, the TTL is increased to the expire entry time value. If the TTL is longer than the expiration timer, the expire entry time value is ignored: no additional time is added to the TTL in this case. Upon expiration, the entry is removed from the DNS lookup table. Removing an entry requires that the table be recompiled, so frequent removals can increase the processing load on the device. Because some DNS entries can have very short TTL (as short as three seconds), you can use this setting to virtually extend the TTL. The default is 1 minute (that is, the minimum TTL for all resolutions is 1 minute). The range is 1 to 65535 minutes.

  Note that for systems running 7.0 or earlier, the expiration time is actually added to the TTL: it does not specify a minimum value.

- **Poll Timer** specifies the time limit after which the device queries the DNS server to resolve the FQDN that was defined in a network object. An FQDN is resolved periodically either when the poll timer has expired, or when the TTL of the resolved IP entry has expired, whichever occurs first.

Step 7 Enable DNS lookups on all interfaces or on specific interfaces. These choices also affect which routing tables are used.

Note that enabling DNS lookups on an interface is not the same as specifying the source interface for lookups. The Firewall Threat Defense always uses a route lookup to determine the source interface. Management-only interfaces other than the dedicated Management interface cannot be used.

- No interfaces selected—Enables DNS lookups on all interfaces. The Firewall Threat Defense checks the data routing table only.
- Specific interfaces selected but not the **Enable DNS Lookup via diagnostic/management interface also** option—Enables DNS lookups on the specified interfaces. The Firewall Threat Defense checks the data routing table only.
- Specific interfaces selected plus the **Enable DNS Lookup via diagnostic/management interface also** option—Enables DNS lookups on the specified interfaces and the Management interface. The Firewall Threat Defense checks the data routing table, and if no route is found, falls back to the management-only routing table.
- Only the **Enable DNS Lookup via diagnostic/management interface also** option—Enables DNS lookups on Management. The Firewall Threat Defense checks only the management-only routing table.

Step 8 To configure the trusted DNS servers, click the **Trusted DNS Servers** tab.

Step 9 By default, the existing DNS servers that are configured in DHCP pool, DHCP relay, DHCP client, or DNS server group are included as trusted DNS servers. If you want to exclude any of them, uncheck the appropriate check boxes.

Step 10 To add trusted DNS servers, under **Specify DNS Servers**, click **Edit**.

Step 11 In the **Select DNS Servers** dialog box, either choose a host object as the trusted DNS server or directly specify the IP address of the trusted DNS server:

a) To choose existing host objects, under **Available Host Objects**, select the required host object and click **Add** to include it to **Selected DNS Servers**. For information on adding the host objects, see Creating Network Objects, on page 1403.

b) To directly provide the IP address(IPv4 or IPv6) of the trusted DNS server, enter the address in the given text field, and click **Add** to include it to **Selected DNS Servers**.

c) Click **Save**. The added DNS servers are displayed in the **Trusted DNS Servers** page.

Note
You can configure a maximum of 12 DNS servers per policy.

Step 12 (Optional) To search for a DNS server that was added, using either the host name or the IP address, use the search field under **Specify DNS Servers**.

Step 13 Click **Save**.

What to do next

To use FQDN objects for access control rules, create an FQDN network object which can then be assigned to an access control rule. For instructions see, Creating Network Objects, on page 1403.

External Authentication

Note You must have administrator privileges to perform this task.

When you enable external authentication for management users, the Firewall Threat Defense verifies the user credentials with an LDAP or RADIUS server as specified in an external authentication object.

Sharing External Authentication Objects

External authentication objects can be used by the Firewall Management Center and Firewall Threat Defense devices. You can share the same object between the Firewall Management Center and devices, or create separate objects. Note that the Firewall Threat Defense supports defining users on the RADIUS server, while the Firewall Management Center requires you to predefine the user list in the external authentication object. You can choose to use the predefined list method for the Firewall Threat Defense, but if you want to define users on the RADIUS server, you must create separate objects for the Firewall Threat Defense and the Firewall Management Center.

Note The timeout range is different for the Firewall Threat Defense and the Firewall Management Center, so if you share an object, be sure not to exceed the Firewall Threat Defense's smaller timeout range (1-30 seconds for LDAP, and 1-300 seconds for RADIUS). If you set the timeout to a higher value, the Firewall Threat Defense external authentication configuration will not work.

Assigning External Authentication Objects to Devices

For the Firewall Management Center, enable the external authentication objects directly on **System** (⚙) > **Users** > **External Authentication**; this setting only affects Firewall Management Center usage, and it does not need to be enabled for managed device usage. For Firewall Threat Defense devices, you must enable the external authentication object in the platform settings that you deploy to the devices, and you can only activate one external authentication object per policy. An LDAP object with CAC authentication enabled cannot also be used for CLI access. Be sure that both the Firewall Threat Defense and the Firewall Management Center can reach the LDAP server, even if you are not sharing the object. The Firewall Management Center is essential to retrieving the user list and downloading it to the device.

Firewall Threat Defense Supported Fields

Only a subset of fields in the external authentication object are used for Firewall Threat Defense SSH access. If you fill in additional fields, they are ignored. If you also use this object for the Firewall Management Center, those fields will be used. This procedure only covers the supported fields for the Firewall Threat Defense. For other fields, see *Configure External Authentication for the Firewall Management Center* in the Cisco Secure Firewall Management Center Administration Guide.

Usernames

Usernames must be Linux-valid usernames and be lower-case only, using alphanumeric characters plus period (.) or hyphen (-). Other special characters such as at sign (@) and slash (/) are not supported. You cannot add the **admin** user for external authentication. You can only add external users (as part of the External Authentication object) in the Firewall Management Center; you cannot add them at the CLI. Note that internal users can only be added at the CLI, not in the Firewall Management Center.

If you previously configured the same username for an internal user using the **configure user add** command, the Firewall Threat Defense first checks the password against the internal user, and if that fails, it checks the AAA server. Note that you cannot later add an internal user with the same name as an external user; only pre-existing internal users are supported. For users defined on the RADIUS server, be sure to set the privilege level to be the same as any internal users; otherwise you cannot log in using the external user password.

Privilege Level

LDAP users always have Config privileges. RADIUS users can be defined as either Config or Basic users.

Before you begin

- SSH access is enabled by default on the management interface. To enable SSH access on data interfaces, see SSH Access, on page 964.
- Inform RADIUS users of the following behavior to set their expectations appropriately:
  - The first time an external user logs in, Firewall Threat Defense creates the required structures but cannot simultaneously create the user session. The user simply needs to authenticate again to start the session. The user will see a message similar to the following: "New external username identified. Please log in again to start a session."
  - If the user's Service-Type attribute is not defined or incorrectly configured in the RADIUS server, and when using the RADIUS-defined users for authentication, the user will see a message similar to the following: "Your username is not defined with a service type that is valid for this system. You are not authorized to access the system?.

    In some cases, the SSH clients close the CLI window on an unsuccessful SSH connection, even before displaying the failure message. Hence, ensure that the user's Service-Type attribute is correctly defined in the RADIUS server.

- Similarly, if the user's Service-Type authorization was changed since the last login, the user will need to re-authenticate. The user will see a message similar to the following: "Your authorization privilege has changed. Please log in again to start a session."

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Click **External Authentication**.

Step 3 Click the **Manage External Authentication Server** link.

You can also open the External Authentication screen by clicking **System** > **Users** > **External Authentication**.

Step 4 Configure an LDAP Authentication Object.

a) Click **Add External Authentication Object**.
b) Set the **Authentication Method** to **LDAP**
c) Enter a **Name** and optional **Description**.
d) Choose a **Server Type** from the drop-down list.
e) For the **Primary Server**, enter a **Host Name/IP Address**.

Note
If you are using a certificate to connect via TLS or SSL, the host name in the certificate must match the host name used in this field. In addition, IPv6 addresses are not supported for encrypted connections.

f) (Optional) Change the **Port** from the default.
g) (Optional) Enter the **Backup Sever** parameters.
h) Enter **LDAP-Specific Parameters**.

- **Base DN**—Enter the base distinguished name for the LDAP directory you want to access. For example, to authenticate names in the Security organization at the Example company, enter `ou=security,dc=example,dc=com`. Alternatively click **Fetch DNs**, and choose the appropriate base distinguished name from the drop-down list.
- (Optional) **Base Filter**—For example, if the user objects in a directory tree have a `physicalDeliveryOfficeName` attribute and users in the New York office have an attribute value of `NewYork` for that attribute, to retrieve only users in the New York office, enter `(physicalDeliveryOfficeName=NewYork)`.
- **User Name**—Enter a distinguished name for a user who has sufficient credentials to browse the LDAP server. For example, if you are connecting to an OpenLDAP server where user objects have a `uid` attribute, and the object for the administrator in the Security division at our example company has a `uid` value of `NetworkAdmin`, you might enter `uid=NetworkAdmin,ou=security,dc=example,dc=com`.
- **Password** and **Confirm Password**—Enter and confirm the password for the user.
- (Optional) **Show Advanced Options**—Configure the following advanced options.
  - **Encryption**—Click **None**, **TLS**, or **SSL**.

Note

If you change the encryption method after specifying a port, you reset the port to the default value for that method. For **None** or **TLS**, the port resets to the default value of 389. If you choose SSL encryption, the port resets to 636.

- **SSL Certificate Upload Path**—For SSL or TLS encryption, you must choose a certificate by clicking **Choose File**.
- (Not Used) **User Name Template**—Not used by the Firewall Threat Defense.
- **Timeout**—Enter the number of seconds before rolling over to the backup connection between 1 and 30. The default is 30.

Note
The timeout range is different for the Firewall Threat Defense and the Firewall Management Center, so if you share an object, be sure not to exceed the Firewall Threat Defense's smaller timeout range (1-30 seconds). If you set the timeout to a higher value, the Firewall Threat Defense external authentication configuration will not work.

i) (Optional) Set the **CLI Access Attribute** if you want to use a shell access attribute other than the user distinguished type. For example, on a Microsoft Active Directory Server, use the `sAMAccountName` shell access attribute to retrieve shell access users by typing `sAMAccountName` in the **CLI Access Attribute** field.

j) Set the **CLI Access Filter**.

Choose one of the following methods:

- To use the same filter you specified when configuring authentication settings, choose **Same as Base Filter**.
- To retrieve administrative user entries based on attribute value, enter the attribute name, a comparison operator, and the attribute value you want to use as a filter, enclosed in parentheses. For example, if all network administrators have a `manager` attribute which has an attribute value of `shell`, you can set a base filter of `(manager=shell)`.

The names on the LDAP server must be Linux-valid usernames:

- Maximum 32 alphanumeric characters, plus hyphen (-) and underscore (_)
- All lowercase
- Cannot start with hyphen (-); cannot be all numbers; cannot include a period (.), at sign (@), or slash (/)

k) Click **Save**.

Step 5 For LDAP, if you later add or delete users on the LDAP server, you must refresh the user list and redeploy the Platform Settings.

a) Choose **System** (⚙) > **Users** > **External Authentication**.

b) Click **Refresh** (↻) next to the LDAP server.

If the user list changed, you will see a message advising you to deploy configuration changes for your device. The Firepower Theat Defense Platform Settings will also show that it is "Out-of-Date on *x* targeted devices."

c) Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Step 6 Configure a RADIUS Authentication Object.

a) Define users on the RADIUS server using the Service-Type attribute.

The following are supported values for the Service-Type attribute:

- Administrator (6)—Provides Config access authorization to the CLI. These users can use all commands in the CLI.
- NAS Prompt (7) or any level other than 6—Provides Basic access authorization to the CLI. These users can use read-only commands, such as **show** commands, for monitoring and troubleshooting purposes.

The names must be Linux-valid usernames:

- Maximum 32 alphanumeric characters, plus hyphen (-) and underscore (_)
- All lowercase
- Cannot start with hyphen (-); cannot be all numbers; cannot include a period (.), at sign (@), or slash (/)

Alternatively, you can predefine users in the external authentication object (see Step 6.k, on page 957). To use the same RADIUS server for the Firewall Threat Defense and Firewall Management Center while using the Service-Type attribute method for the Firewall Threat Defense, create two external authentication objects that identify the same RADIUS server: one object includes the predefined **CLI Access Filter** users (for use with the Firewall Management Center), and the other object leaves the **CLI Access Filter** empty (for use with Firewall Threat Defenses).

b) In Firewall Management Center, click **Add External Authentication Object**.

c) Set the **Authentication Method** to **RADIUS**.

d) Check the **RADIUS Server-Enabled Message Authenticator** check box to require the Message-Authenticator attribute in all RADIUS responses, ensuring that every response from the RADIUS server is securely verified by the Firewall Threat Defense.

This feature is enabled by default for new RADIUS servers. We recommend you enable it for existing servers after the upgrade. Disabling message authenticators may expose your firewalls to potential attacks. Ensure that your RADIUS server has the Message-Authenticator configuration.

e) Enter a **Name** and optional **Description**.

f) For the **Primary Server**, enter a **Host Name/IP Address**.

Note
If you are using a certificate to connect via TLS or SSL, the host name in the certificate must match the host name used in this field. In addition, IPv6 addresses are not supported for encrypted connections.

g) (Optional) Change the **Port** from the default.

h) Enter a **RADIUS Secret Key**.

i) (Optional) Enter the **Backup Sever** parameters.

j) Enter **RADIUS-Specific Parameters**.

- **Timeout (Seconds)**—Enter the number of seconds before rolling over to the backup connection. The default is 30.

- **Retries**—Enter the number of times the primary server connection should be tried before rolling over to the backup connection. The default is 3.

k) (Optional) Instead of using RADIUS-defined users, under **CLI Access Filter**, enter a comma-separated list of usernames in the **Administrator CLI Access User List** field. For example, enter `jchrichton, aerynsun, rygel`.

You may want to use the **CLI Access Filter** method for Firewall Threat Defense so you can use the same external authentication object with Firewall Threat Defense and other platform types. Note that if you want to use RADIUS-defined users, you must leave the **CLI Access Filter** empty.

Make sure that these usernames match usernames on the RADIUS server. The names must be Linux-valid usernames:

- Maximum 32 alphanumeric characters, plus hyphen (-) and underscore (_)
- All lowercase
- Cannot start with hyphen (-); cannot be all numbers; cannot include a period (.), at sign (@), or slash (/)

Note
If you want to only define users on the RADIUS server, you must leave this section empty.

l) Click **Save**.

Step 7 Return to **System** (⚙) > **Users** > **External Authentication**.

Step 8 Click **Refresh** (↻) to view any newly-added objects.

For LDAP when you specify SSL or TLS encryption, you must upload a certificate for the connection; otherwise, the server will not be listed on this window.

Step 9 Click **Slider enabled** (⬭) next to the External Authentication object you want to use. You can only enable one object.

Step 10 Click **Save**.

Step 11 Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Enable Virtual-Router-Aware Interface for External Authentication of Platform

Authentication, Authorization, and Accounting (AAA) for the Firewall Threat Defense device is managed through the management interface of the device. You can also enable virtual-router-aware data interface, data sub-interface, port-channel, or sub port-channel to manage AAA for the Firewall Threat Defense device. When enabled, the AAA route lookup is in the Virtual Routing and Forwarding (VRF) routing domain, and the AAA management traffic is forwarded to the data interfaces. The following server configuration are supported when using virtual-router-aware data interfaces for AAA:

- RADIUS or LDAP servers for external authentication
- FQDN, IPv4, or IPv6 server addresses

To use a virtual-router-aware interface for external authentication of a Firewall Threat Defense device, modify its external authentication policy by associating the authentication servers with the virtual-router-aware interface of the device.

Before you begin

- Ensure that you have configured the required Virtual Routing and Forwarding (VRF) interface with a static route for the device. For information about configuring a VRF interface, see Configure a Virtual Router, on page 1179, and for information about adding a static route, see Add a Static Route, on page 1154.
- Ensure that security zones or interface groups with a single virtual-router-aware interface exists. For information about creating security zones and interface groups, see Create Security Zone and Interface Group Objects, on page 777.
- Ensure that an external authentication server object is configured and enabled. For information on configuring an external authentication object, see Configure External Authentication for the Firewall Threat Defense, on page 217.
- If the primary authentication server is configured with the FQDN of the server, ensure that the backup authentication server, if configured, is also with its FQDN. In addition, configure the DNS server in the threat defense device's management interface. For information about the DNS server configuration, see Configure DNS.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Click **External Authentication** and edit the external authentication policy.

Step 3 In the **External Authentication** dialog box, the available security zone and interface groups are listed. To associate a virtual-router-aware interface with the external authentication servers, select the security zone or interface group having a single virtual-router-aware interface, and then do the following:

a) To associate the interface object with the primary authentication server, click **Add to Primary Server**.

b) (Optional) To associate the interface object with the backup authentication server, click **Add to Backup Server**. If the **Add to Backup Server** button is inactive, it means that a backup server for external authentication is not configured in the device. Configure a backup server from **System** > **Users** > **External Authentication**. For information about configuring a backup authentication server, see Configure External Authentication for the Firewall Threat Defense, on page 217.

Step 4 Click **Ok**.

Step 5 Save and deploy the changes.

Fragment Settings

By default, the Firewall Threat Defense device allows up to 24 fragments per IP packet, and up to 200 fragments awaiting reassembly. You might need to let fragments on your network if you have an application that routinely fragments packets, such as NFS over UDP. However, if you do not have an application that fragments traffic, we recommend that you do not allow fragments by setting **Chain** to 1. Fragmented packets are often used as Denial of Service (DoS) attacks.

Note These settings establish the defaults for devices assigned this policy. You can override these settings for specific interfaces on a device by selecting **Override Default Fragment Setting** in the interface configuration. When you edit an interface, you can find the option on **Advanced** > **Security Configuration**. Select **Devices** > **Device Management**, edit a Firewall Threat Defense device, and select **Interfaces** to edit interface properties..

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Fragment Settings**.

Step 3 Configure the following options. Click **Reset to Defaults** if you want to use the default settings.

- **Size (Block)**—The maximum number of packet fragments from all connections collectively that can be waiting for reassembly. The default is 200 fragments.
- **Chain (Fragment)**—The maximum number of packets into which a full IP packet can be fragmented. The default is 24 packets. Set this option to 1 to disallow fragments.
- **Timeout (Sec)**—The maximum number of seconds to wait for an entire fragmented packet to arrive. The default is 5 seconds. If all fragments are not received within this time, all fragments are discarded.

Step 4 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

HTTP Access

You can enable the HTTPS server to provide a health check mechanism for a cloud load balancer, for example, for the Firewall Threat Defense Virtual on AWS using an Application Load Balancer.

Other uses for HTTPs on the Firewall Threat Defense are not supported; for example, the Firewall Threat Defense does not have a web interface for configuration in this management mode.

This configuration only applies to data interfaces, including any you have configured as management-only. It does not apply to the dedicated Management interface. The Management interface is separate from the other interfaces on the device. It is used to set up and register the device to the Firewall Management Center. It has a separate IP address and static routing.

To use HTTPS, you do not need an access rule allowing the host IP address. You only need to configure HTTPS access according to this section.

You can only use HTTPS to a reachable interface; if your HTTPS host is located on the outside interface, you can only initiate a management connection directly to the outside interface.

Before you begin

- You cannot configure both HTTPS and AnyConnect VPN module of Cisco Secure Client on the same interface for the same TCP port. For example, if you configure remote access SSL VPN on the outside

interface, you cannot also open the outside interface for HTTPS connections on port 443. If you must configure both features on the same interface, use different ports. For example, open HTTPS on port 4443.

- You need network objects that define the hosts or networks you will allow to make HTTPS connections to the device. You can add objects as part of the procedure, but if you want to use object groups to identify a group of IP addresses, ensure that the groups needed in the rules already exist. Select **Objects** > **Object Management** to configure objects.

Note You cannot use the system-provided **any** network object group. Instead, use **any-ipv4** or **any-ipv6**.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **HTTP Access**.

Step 3 Check the **Enable HTTP Server** check box to enable the HTTP server.

Step 4 (Optional) Change the HTTP port. The default is 443.

Step 5 Identify the interfaces and IP addresses that allow HTTP connections.

Use this table to limit which interfaces will accept HTTP connections, and the IP addresses of the clients who are allowed to make those connections. You can use network addresses rather than individual IP addresses.

a) Click **Add** to add a new rule, or click **Edit** to edit an existing rule.

b) Configure the rule properties:

- **IP Address**—The network object or group that identifies the hosts or networks you are allowing to make HTTP connections. Choose an object from the drop-down menu, or click + to add a new network object.
- **Available Zones/Interfaces**—Add the zones that contain the interfaces to which you will allow HTTP connections. For interfaces not in a zone, you can type the interface name into the field below the **Selected Zones/Interfaces** list and click **Add**. These rules will be applied to a device only if the device includes the selected interfaces or zones.

c) Click **OK**.

Step 6 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

ICMP Access

By default, you can send ICMP packets to any interface using either IPv4 or IPv6, with these exceptions:

- The Firewall Threat Defense does not respond to ICMP echo requests directed to a broadcast address.
- The Firewall Threat Defense only responds to ICMP traffic sent to the interface that traffic comes in on; you cannot send ICMP traffic through an interface to a far interface.

To protect the device from attacks, you can use ICMP rules to limit ICMP access to interfaces to particular hosts, networks, or ICMP types. ICMP rules function like access rules, where the rules are ordered, and the first rule that matches a packet defines the action.

If you configure any ICMP rule for an interface, an implicit deny ICMP rule is added to the end of the ICMP rule list, changing the default behavior. Thus, if you want to simply deny a few message types, you must include a permit any rule at the end of the ICMP rule list to allow the remaining message types.

We recommend that you always grant permission for the ICMP unreachable message type (type 3). Denying ICMP unreachable messages disables ICMP path MTU discovery, which can halt IPsec and PPTP traffic. Additionally ICMP packets in IPv6 are used in the IPv6 neighbor discovery process.

Before you begin

Ensure that the objects needed in the rules already exist. Select **Objects** > **Object Management** to configure objects. You need network objects or groups that define the desired hosts or networks, and port objects that define the ICMP message types you want to control.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **ICMP Access**.

Step 3 Configure ICMP rules.

a) Click **Add** to add a new rule, or click **Edit** to edit an existing rule.

b) Configure the rule properties:

- **Action**—Whether to permit (allow) or deny (drop) matching traffic.
- **ICMP Service**—The port object that identifies the ICMP message type.
- **Network**—The network object or group that identifies the hosts or networks whose access you are controlling.
- **Available Zones/Interfaces**—Add the zones that contain the interfaces that you are protecting. For interfaces not in a zone, you can type the interface name into the field below the **Selected Zones/Interfaces** list and click **Add**. These rules will be applied to a device only if the device includes the selected interfaces or zones.

c) Click **OK**.

Step 4 (Optional.) Set rate limits on ICMPv4 Unreachable messages.

- **Rate Limit**—Sets the rate limit of unreachable messages, between 1 and 100 messages per second. The default is 1 message per second.
- **Burst Size**—Sets the burst rate, between 1 and 10. The system sends this number of replies, but subsequent replies are not sent until the rate limit is reached.

Step 5 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

NetFlow

The NetFlow feature enables you to collect IP network traffic information as it enters or exits an interface. The collected traffic information is sent as collected records to a NetFlow Collector server or NetFlow Analyzer. You can analyze the data from NetFlow and determine information, such as source and destination of traffic, class of service, traffic pattern, bandwidth usage, type of traffic, traffic volume, and the causes of the congestion.

With the native NetFlow configuration support, the traffic information collection that was enabled through syslogs flow exports has to be disabled.

The NetFlow provides the option to configure the flow exporter and collectors along with flow event types that must be monitored.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **NetFlow**.

Step 3 Enable the **Enable Flow Export** toggle to enable NetFlow data export.

Step 4 Configure the general NetFlow parameters that controls the frequency of events pushed to the collector.

a) **Active Refresh Interval**—For active connections, specify the time interval (in minutes) between flow-update events.

b) **Delay Flow Create**—Specify the delay (in seconds) before sending a flow-create event. If you do not enter any value, there is no delay and the flow-create event is exported as soon as the flow is created.

c) **Template Timeout Rate**—Specify the time interval (in minutes) at which the template records are sent to the collectors.

Step 5 Click **Add Collector** to configure the collector. See Add Collector in NetFlow, on page 963.

Step 6 Click **Add Traffic Class** to configure the traffic class. See Add Traffic Class to NetFlow, on page 963.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Add Collector in NetFlow

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Choose **NetFlow**.

Step 3 Enable the **Enable Flow Export** toggle to enable NetFlow data export.

Step 4 Click **Add Collector** to configure the collector. You can configure a maximum of five collectors.

Step 5 From the **Host** dropdown list, choose the collector host IP address (IPv4 only) of the NetFlow event collector or server to which the NetFlow packets must be sent. Alternatively, you can click the (+) icon to create a new network host.

Step 6 In the **Port** field, enter the UDP port on the collector to which the NetFlow packets must be sent.

Step 7 From **Available Interfaces or Interface Groups**, choose the interface or interface group through which the collector must be reached. You can choose multiple interfaces or interface groups. The Interface Group object can contain only one interface for a given device. A collector can be reached only through one interface. The object can include a virtual-router-aware interface.

You can click the (+) icon to create a new interface group.

Step 8 Click **Add** to add the interfaces you chose.

Step 9 You can also enter the interface name and click **Add** to add the interface.

Step 10 Click **OK**.

Add Traffic Class to NetFlow

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **NetFlow**.

Step 3 Enable the **Enable Flow Export** toggle to enable NetFlow data export.

Step 4 Click **Add Traffic Class** to configure the traffic class.

Step 5 In the **Name** field, enter the name of the traffic class that must match the NetFlow events.

Step 6 In the **Type** field, choose the traffic class to filter the type of traffic you want to capture:

- **Default**—The traffic class that is matched if none of the traffic classes matches the traffic.
- **Access List**—The specific traffic class that must match the traffic captured for the NetFlow events.

Step 7 If you chose **Access List** as the **Type**, then you must choose the access list object from the **Access List Object** dropdown list.

Note

You can also click the (+) icon to create a new extended access list object. See Configure Extended ACL Objects, on page 1370.

Step 8 In **Event Types**, check the checkboxes for the different NetFlow events that you want to capture and send to the collectors.

Step 9 Click **OK**.

SSH Access

If you enabled Firewall Management Center access on a data interface, such as outside, you should enable SSH on that interface using this procedure. This section describes how to enable SSH connections to one or more *data* interfaces on the Firewall Threat Defense.

The Firewall Threat Defense uses the CiscoSSH stack, which is based on OpenSSH. CiscoSSH supports FIPS compliance and regular updates, including updates from Cisco and the open source community.

Note SSH is enabled by default on the Management interface; however, this screen does not affect Management SSH access.

The Management interface is separate from the other interfaces on the device. It is used to set up and register the device to the Firewall Management Center. SSH for data interfaces shares the internal and external user list with SSH for the Management interface. Other settings are configured separately: for data interfaces, enable SSH and access lists using this screen; SSH traffic for data interfaces uses the regular routing configuration, and not any static routes configured at setup or at the CLI.

For the Management interface, to configure an SSH access list, see the **configure ssh-access-list** command in the Cisco Secure Firewall Threat Defense Command Reference. To configure a static route, see the **configure network static-routes** command. By default, you configure the default route through the Management interface at initial setup.

To use SSH, you do not also need an access rule allowing the host IP address. You only need to configure SSH access according to this section.

You can SSH only to a reachable interface (including an interface in a user-defined virtual router); if your SSH host is located on the outside interface, you can only initiate a management connection directly to the outside interface. When you enable SSH in a user-defined virtual router, and you want VPN users to access SSH, be sure to terminate the VPN on the same virtual router. If the VPN is terminated on another virtual router, then you must configure route leaks between the virtual routers.

SSH supports the following ciphers and key exchange:

- Encryption—aes128-cbc, aes192-cbc, aes256-cbc, aes128-ctr, aes192-ctr, aes256-ctr
- Integrity—hmac-sha2-256
- Key exchange—dh-group14-sha256

Note After you make three consecutive failed attempts to log into the CLI using SSH, the device terminates the SSH connection.

Before you begin

- You can configure SSH internal users at the CLI using the **configure user add** command; see Add an Internal User at the CLI, on page 215. By default, there is an **admin** user for which you configured the password during initial setup. You can also configure external users on LDAP or RADIUS by configuring **External Authentication** in platform settings. See External Authentication, on page 952.
- You need network objects that define the hosts or networks you will allow to make SSH connections to the device. You can add objects as part of the procedure, but if you want to use object groups to identify a group of IP addresses, ensure that the groups needed in the rules already exist. Select **Objects** > **Object Management** to configure objects.

Note You cannot use the system-provided **any** network object. Instead, use **any-ipv4** or **any-ipv6**.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **SSH Access**.

Step 3 Identify the interfaces and IP addresses that allow SSH connections.

Use this table to limit which interfaces will accept SSH connections, and the IP addresses of the clients who are allowed to make those connections. You can use network addresses rather than individual IP addresses.

a) Click **Add** to add a new rule, or click **Edit** to edit an existing rule.

b) Configure the rule properties:

- **IP Address**—The network object or group that identifies the hosts or networks you are allowing to make SSH connections. Choose an object from the drop-down menu, or click + to add a new network object.
- **Available Zones/Interfaces**—Add the zones that contain the interfaces to which you will allow SSH connections. For interfaces not in a zone, you can type the interface name into the field below the **Selected Zones/Interfaces** list and click **Add**. You can also add loopback interfaces and virtual-router-aware interfaces. These rules will be applied to a device only if the device includes the selected interfaces or zones.

c) Click **OK**.

Step 4 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

SMTP Server

You must identity an SMTP server if you configure email alerts in the Syslog settings. The source email address you configure for Syslog must be a valid account on the SMTP servers.

Before you begin

Ensure that the network objects that define the host address of the primary and secondary SMTP servers exist. Select **Objects** > **Object Management** to define the objects. Alternatively, you can create the objects while editing the policy.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Click **SMTP Server**.

Step 3 Select the network objects that identify the **Primary Server IP Address** and optionally, the **Secondary Server IP Address**.

Step 4 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

SNMP

Simple Network Management Protocol (SNMP) defines a standard way for network management stations running on PCs or workstations to monitor the health and status of many types of devices, including switches, routers, and security appliances. You can use the SNMP page to configure a firewall device for monitoring by SNMP management stations.

The Simple Network Management Protocol (SNMP) enables monitoring of network devices from a central location. Cisco security appliances support network monitoring using SNMP versions 1, 2c, and 3, as well as traps and SNMP read access; SNMP write access is not supported.

SNMPv3 supports read-only users and encryption with DES (deprecated), 3DES, AES256, AES192, and AES128.

Note The DES option has been deprecated. If your deployment includes SNMP v3 users using DES encryption that were created using a version previous to 6.5, you can continue to use those users for Firewall Threat Defenses running versions 6.6 and previous. However, you cannot edit those users and retain DES encryption, or create new users with DES encryption. If your Firewall Management Center manages any Firewall Threat Defenses running Versions 7.0+, deploying a platform settings policy that uses DES encryption to those Firewall Threat Defenses will fail.

Note SNMP configuration supports Routed and Diagnostic interface only.

Note To create an alert to an external SNMP server, access **Policies** > **Actions** > **Alerts**

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **SNMP**.

Step 3 Enable SNMP and configure basic options.

- **Enable SNMP Servers**—Whether to provide SNMP information to the configured SNMP hosts. You can deselect this option to disable SNMP monitoring while retaining the configuration information.
- **Read Community String**, **Confirm**—Enter the password used by a SNMP management station when sending requests to the Firewall Threat Defense device. The SNMP community string is a shared secret among the SNMP management stations and the network nodes being managed. The security device uses the password to determine if the incoming SNMP request is valid. The password is a case-sensitive alphanumeric string of up to 32 characters; spaces and special characters are not permitted.
- **System Administrator Name**—Enter the name of the device administrator or other contact person. This string is case-sensitive and can be up to 127 characters. Spaces are accepted, but multiple spaces are shortened to a single space.
- **Location**—Enter the location of this security device (for example, Building 42,Sector 54). This string is case-sensitive and can be up to 127 characters. Spaces are accepted, but multiple spaces are shortened to a single space.
- **Port**—Enter the UDP port on which incoming requests will be accepted. The default is 161.

Step 4 (SNMPv3 only.) Add SNMPv3 Users, on page 973.

Step 5 Add SNMP Hosts, on page 976.

Step 6 Configure SNMP Traps, on page 978.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

About SNMP

SNMP is an application-layer protocol that facilitates the exchange of management information between network devices and is part of the TCP/IP protocol suite. Firewall Threat Defense provides support for network monitoring using SNMP Versions 1, 2c, and 3, and support the use of all three versions simultaneously. The SNMP agent running on the Firewall Threat Defense interface lets you monitor the network devices through network management systems (NMSes), such as HP OpenView. Firewall Threat Defense supports SNMP read-only access through issuance of a GET request. SNMP write access is not allowed, so you cannot make changes with SNMP. In addition, the SNMP SET request is not supported.

You can configure the Firewall Threat Defense to send traps, which are unsolicited messages from the managed device to the management station for certain events (event notifications) to an NMS, or you can use the NMS to browse the Management Information Bases (MIBs) on the security devices. MIBs are a collection of definitions, and the Firewall Threat Defense maintain a database of values for each definition. Browsing a MIB means issuing a series of GET-NEXT or GET-BULK requests of the MIB tree from the NMS to determine values.

An SNMP agent notifies the designated management stations if events occur that are predefined to require a notification, for example, when a link in the network goes up or down. The notification it sends includes an SNMP OID, which identifies itself to the management stations. The agent also replies when a management station asks for information.

SNMP Terminology

The following table lists the terms that are commonly used when working with SNMP.

Table 55: SNMP Terminology

Term	Description
Agent	The SNMP server running on the Secure Firewall Threat Defense. The SNMP agent has the following features: • Responds to requests for information and actions from the network management station. • Controls access to its Management Information Base, the collection of objects that the SNMP manager can view or change. • Does not allow SET operations.
Browsing	Monitoring the health of a device from the network management station by polling required information from the SNMP agent on the device. This activity may include issuing a series of GET-NEXT or GET-BULK requests of the MIB tree from the network management station to determine values.
Management Information Bases (MIBs)	Standardized data structures for collecting information about packets, connections, buffers, failovers, and so on. MIBs are defined by the product, protocols, and hardware standards used by most network devices. SNMP network management stations can browse MIBs and request specific data or events be sent as they occur.
Network management stations (NMSs)	The PCs or workstations set up to monitor SNMP events and manage devices.
Object identifier (OID)	The system that identifies a device to its NMS and indicates to users the source of information monitored and displayed.

Term	Description
Trap	Predefined events that generate a message from the SNMP agent to the NMS. Events include alarm conditions such as linkup, linkdown, coldstart, warmstart, authentication, or syslog messages. **coldStart**—The coldStart trap occurs when the SNMP agent starts after the SNMP configuration. This trap also occurs when the agent starts after a system reboot. **Note** For cluster and HA nodes, post a reload, if the interfaces reboot time exceeds 5 minutes (preset threshold), the trap is dropped. When the cluster and HA nodes have rebooted sucessfully, all other traps are sent as expected. The **snmp-server enable traps snmp coldstart** command is used to enable and disable transmission of these traps.

MIBs and Traps

MIBs are either standard or enterprise-specific. Standard MIBs are created by the IETF and documented in various RFCs. A trap reports significant events occurring on a network device, most often errors or failures. SNMP traps are defined in either standard or enterprise-specific MIBs. Standard traps are created by the IETF and documented in various RFCs. SNMP traps are compiled into the ASA software.

If needed, you can also download RFCs, standard MIBs, and standard traps from the following locations:

http://www.ietf.org/

Browse the SNMP Object Navigator to look up Cisco MIBs, traps, and OIDs from the following location:

https://snmp.cloudapps.cisco.com/Support/SNMP/do/BrowseOID.do?local=en

In addition, download Cisco OIDs by FTP from the following location:

ftp://ftp.cisco.com/pub/mibs/oid/oid.tar.gz

Supported Tables and Objects in MIBs

The following sections list the supported tables and objects for the specified MIBs.

Remote Access VPN Polling

Table 56: CISCO-REMOTE-ACCESS-MONITOR-MIB

Counter	OID	Description
Active Sessions	crasNumSessions (1.3.6.1.4.1.9.9.392.1.3.1)	The number of currently active sessions.
Users	crasNumUsers (1.3.6.1.4.1.9.9.392.1.3.3)	The number of users who have active sessions.
Peak Sessions	crasNumPeakSessions (1.3.6.1.4.1.9.9.392.1.3.41)	The number of peak RA sessions since system up.

Site-to-Site VPN Tunnel Polling

Table 57: CISCO-REMOTE-ACCESS-MONITOR-MIB

Counter	OID	Description
LAN to LAN Sessions	crasL2LNumSessions (1.3.6.1.4.1.9.9.392.1.3.29)	The number of currently active LAN to LAN sessions.
Peak LAN to LAN Sessions	crasL2LPeakConcurrentSessions (1.3.6.1.4.1.9.9.392.1.3.31)	The number of peak concurrent LAN to LAN sessions since the system is up.

Connection Polling

Table 58: CISCO-FIREWALL-MIB

Counter	OID	Description
Active Connections	cfwConnectionActive (1.3.6.1.4.1.9.9.147.1.2.2.2.1.3.40.6)	The number of connections currently in use by the entire firewall.
Peak Connections	cfwConnectionPeak (1.3.6.1.4.1.9.9.147.1.2.2.2.1.3.40.7)	The highest number of connections in use at any one time since system startup.
Connections Per Second	cfwConnectionPerSecond (1.3.6.1.4.1.9.9.147.1.2.2.3)	The current connections per second rate on the firewall.
Peak Connections Per Second	cfwConnectionPerSecondPeak (1.3.6.1.4.1.9.9.147.1.2.2.4)	The highest number of connections per second on the firewall since system startup.

NAT Translation Polling

Table 59: CISCO-NAT-EXT-MIB

Counter	OID	Description
Active Translations	cneAddrTranslationNumActive (1.3.6.1.4.1.9.9.532.1.1.1.1)	The total number of address translation entries that are currently available in the NAT device. This indicates the aggregate of the translation entries created from both the static and dynamic address translation mechanisms.
Peak Active Translations	cneAddrTranslationNumPeak (1.3.6.1.4.1.9.9.532.1.1.1.2)	The maximum number of address translation entries that are active at any one time since the system startup. This indicates the high watermark of address translation entries that are active at any one time since the system startup. This object includes the translation entries created from both the static and dynamic address translation mechanisms.

Routing Table Entries Polling

Table 60: IP-FORWARD-MIB

Counter	OID	Description
Active Translations	inetCidrRouteNumber (1.3.6.1.2.1.4.24.6)	The total number of current inetCidrRouteTable entries that valid.

Interface Duplex Status Polling

Table 61: CISCO-IF-EXTENSION-MIB

Counter	OID	Description
Duplex Status	cieIfDuplexCfgStatus (1.3.6.1.4.1.9.9.276.1.1.2.1.20)	This object specifies the configured duplex status on the given interface.
Detected Duplex Status	cieIfDuplexDetectStatus (1.3.6.1.4.1.9.9.276.1.1.2.1.21)	This object specifies the detected duplex status on the given interface.

Snort 3 Intrusion Event Rate Polling

Table 62: CISCO-UNIFIED-FIREWALL-MIB

Counter	OID	Description
Snort 3 Intrusion Event Rate	cufwAaicIntrusionEvtRate (1.3.6.1.4.1.9.9.491.1.5.3.2.1)	The rate at which intrusion events were recorded by Snort on this firewall averaged over the last 300 seconds.

BGP Peer-Flap Trap Notification

Table 63: BGP4-MIB

Counter	OID	Description
BGP Peer-flap	bgpBackwardTransition (1.3.6.1.4.1.9.9.491.1.5.3.2.1)	The BGPBackwardTransition Event is generated when the BGP FSM moves from a higher numbered state to a lower numbered state.

CPU Utilization Polling

Table 64: CISCO-PROCESS-MIB

Counter	OID	Description
CPU Total Utilization	cpmCPUTotal1minRev (1.3.6.1.4.1.9.9.109.1.1.1.1.7.1)	Total system CPU utilization for the last one minute.

Counter	OID	Description
Individual CPU Core Utilization	Associated parameters and values of cpmCPUTotal1minRev 1.3.6.1.4.1.9.9.109.1.1.1.1.7.2 to 1.3.6.1.4.1.9.9.109.1.1.1.1.7.(n+1)	Individual CPU core utilization values for the last one minute, where 'n' represents the number of cores. Examples: • 3.6.1.4.1.99.109.1.1.1.1.7.(n+2) - Aggregate system CPU utilization % (This value is same as the system cpu usage from .3.6.1.4.1.99.109.1.1.1.1.7.1 in single context mode). • 3.6.1.4.1.99.109.1.1.1.1.7.(n+3) - Snort average CPU utilization % (total aggregate value of all snort instances) • 3.6.1.4.1.99.109.1.1.1.1.7.(n+4) - System process average % (average of "Sysproc" cores)

Note The SNMP OIDs 1.3.6.1.2.1.25.3.3 and 1.3.6.1.2.1.25.3.4 pertaining to CPU monitoring (hrProcessorTable and hrNetworkTable) were removed on ASA FirePOWER. You can view and monitor the CPU health details of the device only through its device manager.

In ENTITY-MIB, two new vendor OIDs were added for the dual EPM 2X100G and 4X200G cards of Secure Firewall 4200 for slot 2 and slot 3—cevFPRNM4X200Gng and cevFPRNM2X100Gng.

Add SNMPv3 Users

Note You create users for SNMPv3 only. These steps are not applicable for SNMPv1 or SNMPv2c.

Note that SNMPv3 only supports read-only users.

SNMP users have a specified username, an authentication password, an encryption password, and authentication and encryption algorithms to use.

Note When using SNMPv3 with clustering or High Availability, if you add a new cluster unit after the initial cluster formation or you replace a High Availability unit, then SNMPv3 users are not replicated to the new unit. You must remove the users, re-add them, and then redeploy your configuration to force the users to replicate to the new unit.

Authentication algorithm options

For a successful SNMPv3 user authentication, the authentication key size should be equal to or larger than the encryption key size. For example, a user authentication using SHA-1 with AES256 fails, whereas authentication using SHA256 with AES256 succeeds.

Versions	Supported Algorithm		Deprecated Algorithm	
	Authentication	Encryption	Authentication	Encryption
< 6.4	MD5 and SHA	DES, 3DES, AES256, AES192, and AES128	-	-
6.5	SHA	3DES, AES256, AES192, and AES128	MD5*	DES**
6.6 and 6.7	SHA, and SHA256	3DES, AES256, AES192, and AES128	-	-
7.0+	SHA, SHA224, SHA256, and SHA384	3DES, AES256, AES192, and AES128	-	-

* If your deployment includes SNMP v3 users using the MD5 authentication algorithm that were created using a version previous to 6.5, you can continue to use those users for Firewall Threat Defense devices running versions 6.7 and previous. However, you cannot edit those users and retain the MD5 authentication algorithm, or create new users with the MD5 authentication algorithm. If your Firewall Management Center manages any Firewall Threat Defensedevices running versions 7.0+, deploying a platform settings policy that uses the MD5 authentication algorithm to those devices will fail.

** If your deployment includes SNMP v3 users using DES encryption that were created using a version previous to 6.5, you can continue to use those users for Firewall Threat Defensedevices running versions 6.7 and previous. However, you cannot edit those users and retain DES encryption, or create new users with DES encryption. If your Firewall Management Center manages any Firewall Threat Defense devices running versions 7.0+, deploying a platform settings policy that uses DES encryption to those devices will fail.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Click **SNMP** > **Users**.

Step 3 Click **Add**.

Step 4 Select the security level for the user from the **Security Level** drop-down list.

- **Auth**—Authentication but No Privacy, which means that messages are authenticated.
- **No Auth**—No Authentication and No Privacy, which means that no security is applied to messages.
- **Priv**—Authentication and Privacy, which means that messages are authenticated and encrypted.

Step 5 Enter the name of the SNMP user in the **Username** field. Usernames must be 32 characters or less.

Step 6 Select the type of password, you want to use in the **Encryption Password Type** drop-down list.

- **Clear text**—The Firewall Threat Defense device will still encrypt the password when deploying to the device.
- **Encrypted**—The Firewall Threat Defense device will directly deploy the encrypted password.

Step 7 In the **Auth Algorithm Type** drop-down list, select the type of authentication you want to use: SHA, SHA224, SHA256, or SHA384.

Note
The MD5 option has been deprecated. If your deployment includes SNMP v3 users using the MD5 authentication algorithm that were created using a version previous to 6.5, you can continue to use those users for FTDs running versions 6.7 and previous. However, you cannot edit those users and retain the MD5 authentication algorithm, or create new users with the MD5 authentication algorithm. If your Firewall Management Center manages any Firewall Threat Defenses running Versions 7.0+, deploying a platform settings policy that uses the MD5 authentication algorithm to those Firewall Threat Defenses will fail.

Step 8 In the **Authentication Password** field, enter the password to use for authentication. If you selected Encrypted as the Encrypt Password Type, the password must be formatted as xx:xx:xx..., where xx are hexadecimal values.

Note
The length of the password will depend on the authentication algorithm selected. For all passwords, the length must be 256 characters or less.

If you selected Clear Text as the Encrypt Password Type, repeat the password in the **Confirm** field.

Step 9 In the **Encryption Type** drop-down list, select the type of encryption you want to use: AES128, AES192, AES256, 3DES.

Note
To use AES or 3DES encryption, you must have the appropriate license installed on the device.

Note
The DES option has been deprecated. If your deployment includes SNMP v3 users using DES encryption that were created using a version previous to 6.5, you can continue to use those users for Firewall Threat Defenses running versions 6.7 and previous. However, you cannot edit those users and retain DES encryption, or create new users with DES encryption. If your Firewall Management Center manages any Firewall Threat Defenses running Versions 7.0+, deploying a platform settings policy that uses DES encryption to those Firewall Threat Defenses will fail.

Step 10 Enter the password to use for encryption in the **Encryption Password** field. If you selected Encrypted as the Encrypt Password Type, the password must be formatted as xx:xx:xx..., where xx are hexadecimal values. For encrypted passwords, the length of the password depends on the encryption type selected. The password sizes are as follows (where each xx is one octal):

- AES 128 requires 16 octals
- AES 192 requires 24 octals
- AES 256 requires 32 octals
- 3DES requires 32 octals
- DES can be any size

Note
For all passwords, the length must be 256 characters or less.

If you selected Clear Text as the Encrypt Password Type, repeat the password in the **Confirm** field.

Step 11 Click **OK**.

Step 12 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Add SNMP Hosts

Use Host to add or edit entries in the SNMP Hosts table on the SNMP page. These entries represent SNMP management stations allowed to access the Firewall Threat Defense device.

You can add up to 8192 hosts. However, only 128 of this number can be for traps.

Note In 7.4 and later, the Management and Diagnostic interfaces are merged. If Platform Settings for syslog servers or SNMP hosts specify the Diagnostic interface by name, then you must use separate Platform Settings policies for merged and non-merged devices (7.3 and earlier, and some upgraded 7.4 FTDs).

Before you begin

Ensure that the network objects that define the SNMP management stations exist. Select **Device** > **Object Management** to configure network objects.

Note The supported network objects include IPv6 hosts, IPv4 hosts, IPv4 range and IPv4 subnet addresses.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Click **SNMP** > **Hosts**.

Step 3 Click **Add**.

Step 4 In the **IP Address** field, either enter a valid IPv6 or IPv4 host or select the network object that defines the SNMP management station's host address.

The IP address can be an IPv6 host, IPv4 host, IPv4 range or IPv4 subnet.

Step 5 Select the appropriate SNMP version from the **SNMP version** drop-down list.

Step 6 (SNMPv3 only.) Select the username of the SNMP user that you configured from the **User Name** drop-down list.

Note
You can associate up to 23 SNMP users per SNMP host.

Step 7 (SNMPv1, 2c only.) In the **Read Community String** field, enter the community string that you have already configured, for read access to the device. Re-enter the string to confirm it.

Note
This string is required, only if the string used with this SNMP station is different from the one already defined in the **Enable SNMP Server** section.

Step 8 Select the type of communication between the device and the SNMP management station. You can select both types.

- **Poll**—The management station periodically requests information from the device.
- **Trap**—The device sends trap events to the management station as they occur.

Note
When the SNMP host IP address is either an IPv4 range or an IPv4 subnet, you can configure either **Poll** or **Trap**, not both.

Step 9 In the **Port** field, enter a UDP port number for the SNMP host. The default value is 162. The valid range is 1 to 65535.

Step 10 Select the interface type for communication between the device and the SNMP management station under the **Reachable By** options. You can select either the device's Management interface or an available security zone/named interface.

- **Device Management Interface**—Communication between the device and the SNMP management station occurs over the Management interface.
  - When you choose this interface for SNMPv3 polling, all configured SNMPv3 users are allowed to poll and are not restricted to the user chosen in Step 6, on page 977. Here, SNMPv1 and SNMPv2c are not allowed from an SNMPv3 host.
  - When you choose this interface for SNMPv1 and SNMPv2c polling, the polling is not restricted at all to the version selected in Step 5, on page 977.
- **Security Zones or Named Interface**—Communication between the device and the SNMP management station occurs over a security zone or interface.
  - Search for zones in the **Available Zones** field.
  - Add the zones that contain the interfaces through which the device communicates with the management station to the **Selected Zone/Interface** field. For interfaces not in a zone, you can type the interface name into the field below the **Selected Zone/Interface** list and click **Add**. You can also choose a loopback interface and virtual-router-aware interfaces. The host will be configured on a device only if the device includes the selected interfaces or zones.

Step 11 Click **OK**.

Step 12 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure SNMP Traps

Use SNMP Traps to configure SNMP traps (event notifications) for the Firewall Threat Defense device. Traps are different from browsing; they are unsolicited "comments" from the Firewall Threat Defense device to the management station for certain events, such as linkup, linkdown, and syslog event generated. An SNMP object ID (OID) for the device appears in SNMP event traps sent from the device.

Some traps are not applicable to certain hardware models. These traps will be ignored if you apply the policy to one of these models. For example, not all models have field-replaceable units, so the **Field Replaceable Unit Insert/Delete** trap will not be configured on those models.

SNMP traps are defined in either standard or enterprise-specific MIBs. Standard traps are created by the IETF and documented in various RFCs. SNMP traps are compiled into the Firewall Threat Defense software.

If needed, you can download RFCs, standard MIBs, and standard traps from the following location:

http://www.ietf.org/

Browse the complete list of Cisco MIBs, traps, and OIDs from the following location:

SNMP Object Navigator

In addition, download Cisco OIDs by FTP from the following location:

ftp://ftp.cisco.com/pub/mibs/oid/oid.tar.gz

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Click **SNMP** > **SNMP Traps** to configure SNMP traps (event notifications) for the Firewall Threat Defense device.

Step 3 Select the appropriate Enable Traps options. You can select either or both options.

a) Check **Enable All SNMP Traps** to quickly select all traps in the subsequent four sections.

b) Check **Enable All Syslog Traps** to enable transmission of trap-related syslog messages.

Note
SNMP traps are of higher priority than other notification messages from the Firewall Threat Defense as they are expected to be near real-time. When you enable all SNMP or syslog traps, it is possible for the SNMP process to consume excess resources in the agent and in the network, causing the system to hang. If you notice system delays, unfinished requests, or timeouts, you can selectively enable SNMP and syslog traps. You can also limit the rate at which syslog messages are generated by severity level or message ID. For example, all syslog message IDs that begin with the digits 212 are associated with the SNMP class; see Limit the Rate of Syslog Message Generation, on page 996.

Step 4 The event-notification traps in the **Standard** section are enabled by default for an existing policy:

- **Authentication** – Unauthorized SNMP access. This authentication failure occurs for packets with an incorrect community string.
- **Link Up** – One of the device's communication links has become available (it has "come up"), as indicated in the notification.
- **Link Down** – One of the device's communication links has failed, as indicated in the notification.
- **Cold Start** – The device is reinitializing itself such that its configuration or the protocol entity implementation may be altered.
- **Warm Start** – The device is reinitializing itself such that its configuration and the protocol entity implementation is unaltered.

Step 5 Select the desired event-notification traps in the **Entity MIB** section:

- **Field Replaceable Unit Insert** – A Field Replaceable Unit (FRU) has been inserted, as indicated. (FRUs include assemblies such as power supplies, fans, processor modules, interface modules, etc.)
- **Field Replaceable Unit Delete** – A Field Replaceable Unit (FRU) has been removed, as indicated in the notification
- **Configuration Change** – There has been a hardware change, as indicated in the notification

Step 6 Select the desired event-notification traps in the **Resource** section:

- **Connection Limit Reached** – This trap indicates that a connection attempt was rejected because the configured connections limit has been reached.

Step 7 Select the desired event-notification traps in the **Other** section:

- **NAT Packet Discard** – This notification is generated when IP packets are discarded by the NAT function. Available Network Address Translation addresses or ports have fallen below configured threshold.
- **CPU Rising Threshold** – This notification is generated when rising CPU utilization exceeds a predefined threshold for a configured period of time. Check this option to enable CPU rising threshold notifications:
  - **Percentage** – The default value is 70 percent for the high threshold notification; the range is between 10 and 94 percent. The critical threshold is hardcoded at 95 percent.
  - **Period** – The default monitoring period is 1 minute; the range is between 1 and 60 minutes.
- **Memory Rising Threshold** – This notification is generated when rising memory utilization exceeds a predefined threshold, thus reducing available memory. Check this option to enable memory rising threshold notifications:
  - **Percentage** – The default value is 70 percent for the high threshold notification; the range is between 50 and 95 percent.
- **Failover** – This notification is generated when there is a change in the failover state as reported by the CISCO-UNIFIED-FIREWALL-MIB.
- **Cluster** – This notification is generated when there is a change in the cluster health as reported by the CISCO-UNIFIED-FIREWALL-MIB.
- **Peer Flap** – This notification is generated when there is BGP route flapping, a situation in which BGP systems send an excessive number of update messages to advertise network reachability information.

Step 8 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

SSL

Note You must have administrator privileges and be in a leaf domain to perform this task.

You must make sure that you are running a fully licensed version of the Secure Firewall Management Center. The SSL Settings will be disabled if you are running Secure Firewall Management Center in evaluation mode. Additionally, the SSL Settings will be disabled when the licensed Secure Firewall Management Center version does not meet the export-compliance criteria. If you are using Remote Access VPN with SSL, your Smart Account must have the strong-crypto features enabled. For more information, see *License Types and Restrictions* in the Cisco Secure Firewall Management Center Administration Guide.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit a Firewall Threat Defense policy.

Step 2 Select **SSL**.

Step 3 Add entries to the **Add SSL Configuration** table.

a) Click **Add** to create a new entry, or click **Edit** if the entry already exists.

b) Select the required security configurations from the drop-down list .

- **Protocol Version**—Specifies the TLS protocols to be used while establishing remote access VPN sessions.
- **Security Level**—Indicates the kind of security positioning you would like to set up for the SSL.

Step 4 Select the **Available Algorithms** based on the protocol version that you select and click **Add** to include them for the selected protocol. For more information, see About SSL Settings, on page 980.

The algorithms are listed based on the protocol version that you select. Each security protocol identifies unique algorithm for setting up the security level.

Step 5 Click **OK** to save the changes.

What to do next

Select **Deploy** > **Deployment** and click **Deploy** to deploy the policy to the assigned devices.

About SSL Settings

The Firewall Threat Defense device uses the Secure Sockets Layer (SSL) protocol and Transport Layer Security (TLS) to support secure message transmission for Remote Access VPN connection from remote

clients. The SSL Settings window lets you configure SSL versions and encryption algorithms that will be negotiated and used for message transmission during remote VPN access over SSL.

Note Though you configure Firewall Management Center and Firewall Threat Defense to operate in a security certifications (UCAPL, CC, or FIPS) compliance mode, the Firewall Management Center allows configuration of unsupported ciphers. For example, in a FIPS enabled mode, management center allows configuring DH Group 5 which is not FIPS-compliant. However, VPN tunnel does not negotiate due to the non-compliant cipher usage.

Configure the SSL Settings at the following location:

Devices > **Platform Settings**, and then click **SSL**.

Fields

Minimum SSL Version as Server—Specify the minimum SSL/TLS protocol version that the Firewall Threat Defense device uses when acting as a server. For example, when it functions as a Remote Access VPN Gateway.

TLS Version—Select one of the following TLS versions from the drop-down list:

TLS V1	Accepts SSLv2 client hellos and negotiates TLSv1 (or greater).
TLSV1.1	Accepts SSLv2 client hellos and negotiates TLSv1.1 (or greater).
TLSV1.2	Accepts SSLv2 client hellos and negotiates TLSv1.2 (or greater).
TLSV1.3	Accepts SSLv2 client hellos and negotiates TLSv1.3 (or greater).

Note TLS 1.3 in remote access VPN requires Cisco Secure Client, Version 5.0 and above.

DTLS Version—Select the DTLS versions from the drop-down list, based on the selected TLS version. By default, DTLSv1 is configured on Firewall Threat Defense devices, you can choose the DTLS version as per your requirement.

Note Ensure that the TLS protocol version is higher than or equal to the DTLS protocol version selected. TLS protocol versions support the following DTLS versions:

TLS V1	DTLSv1
TLSV1.1	DTLSv1
TLSV1.2	DTLSv1, DTLSv1.2
TLSV1.3	DTLSv1, DTLSv1.2

Diffie-Hellman Group—Choose a group from the drop-down list. Available options are Group1 - 768-bit modulus, Group2 - 1024-bit modulus, Group5 - 1536-bit modulus, Group14 - 2048-bit modulus, 224-bit prime order, and Group24 - 2048-bit modulus, 256-bit prime order. The default is Group1.

Elliptical Curve Diffie-Hellman Group—Choose a group from the drop-down list. Available options are Group19 - 256-bit EC, Group20 - 384-bit EC, and Group21 - 521-bit EC. The default value is Group19.

TLSv1.2 adds support for the following ciphers:

- ECDHE-ECDSA-AES256-GCM-SHA384
- ECDHE-RSA-AES256-GCM-SHA384
- DHE-RSA-AES256-GCM-SHA384
- AES256-GCM-SHA384
- ECDHE-ECDSA-AES256-SHA384
- ECDHE-RSA-AES256-SHA384
- ECDHE-ECDSA-AES128-GCM-SHA256
- ECDHE-RSA-AES128-GCM-SHA256
- DHE-RSA-AES128-GCM-SHA256
- RSA-AES128-GCM-SHA256
- ECDHE-ECDSA-AES128-SHA256
- ECDHE-RSA-AES128-SHA256

Note ECDSA and DHE ciphers are the highest priority.

TLSv1.3 adds support for the following ciphers:

- TLS_AES_128_GCM_SHA256
- TLS_CHACHA20_POLY1305_SHA256
- TLS_AES_256_GCM_SHA384

The SSL configuration table can be used to specify the protocol version, security level, and Cipher algorithms that you want to support on the Secure Firewall Threat Defense devices.

Protocol Version—Lists the protocol version that the Secure Firewall Threat Defense device supports and uses for SSL connections. Available protocol versions are:

- Default
- TLSV1
- TLSV1.1
- TLSV1.2
- TLSV1.3

- DTLSv1
- DTLSv1.2

Security Level—Lists the cipher security levels that Firewall Threat Defense device supports and uses for SSL connections.

If you have Firewall Threat Defense devices with evaluation license, the security level is Low by default. With Firewall Threat Defense smart license, the default security level is High. You can choose one of the following options to configure the required security level:

- **All** includes all ciphers, including NULL-SHA.
- **Low** includes all ciphers, except NULL-SHA.
- **Medium** includes all ciphers, except NULL-SHA, DES-CBC-SHA, RC4-SHA, and RC4-MD5 (this is the default).
- **Fips** includes all FIPS-compliant ciphers, except NULL-SHA, DES-CBC-SHA, RC4-MD5, RC4-SHA, and DES-CBC3-SHA, TLS_CHACHA20_POLY1305_SHA256.
- **High** includes only AES-256 with SHA-2 ciphers and applies to TLS version 1.2 and the *default* version.
- **Custom** includes one or more ciphers that you specify in the Cipher algorithms/custom string box. This option provides you with full control of the cipher suite using OpenSSL cipher definition strings.

Cipher Algorithms/Custom String—Lists the cipher algorithms that the Firewall Threat Defense device supports and uses for SSL connections. For more information about ciphers using OpenSSL, see https://www.openssl.org/docs/apps/ciphers.html

Ciphers supported by TLSv1.2 only

ECDHE-ECDSA-AES256-GCM-SHA384
ECDHE-RSA-AES256-GCM-SHA384
DHE-RSA-AES256-GCM-SHA384
AES256-GCM-SHA384
ECDHE-ECDSA-AES256-SHA384
ECDHE-RSA-AES256-SHA384
DHE-RSA-AES256-SHA256
AES256-SHA256
ECDHE-ECDSA-AES128-GCM-SHA256
ECDHE-RSA-AES128-GCM-SHA256
DHE-RSA-AES128-GCM-SHA256
AES128-GCM-SHA256
ECDHE-ECDSA-AES128-SHA256

ECDHE-RSA-AES128-SHA256
DHE-RSA-AES128-SHA256
AES128-SHA256

Note ECDSA and DHE ciphers are the highest priority.

Ciphers not supported by TLSv1.1 or TLSv1.2

RC4-SHA
RC4-MD5
DES-CBC-SHA
NULL-SHA

Syslog

You can enable system logging (syslog) for Firewall Threat Defense devices. Logging information can help you identify and isolate network or device configuration problems. You can also send some security events to a syslog server. The following topics explain logging and how to configure it.

About Syslog

System logging is a method of collecting messages from devices to a server running a syslog daemon. Logging to a central syslog server helps in aggregation of logs and alerts. Cisco devices can send their log messages to a UNIX-style syslog service. A syslog service accepts messages and stores them in files, or prints them according to a simple configuration file. This form of logging provides protected long-term storage for logs. Logs are useful both in routine troubleshooting and in incident handling.

Table 65: System Logs for Secure Firewall Threat Defense

Logs Related To	Details	Configure In
Device and system health, network configuration	This syslog configuration generates messages for features running on the data plane, that is, features that are defined in the CLI configuration that you can view with the **show running-config** command. This includes features such as routing, VPN, data interfaces, DHCP server, NAT, and so forth. Data plane syslog messages are numbered, and they are the same as those generated by devices running ASA software. However, Secure Firewall Threat Defense does not necessarily generate every message type that is available for ASA Software. For information on these messages, see *Cisco Secure Firewall Threat Defense Syslog Messages* at https://www.cisco.com/c/en/us/td/docs/security/firepower/Syslogs/b_fptd_syslog_guide.html. This configuration is explained in the following topics.	**Platform Settings**

Logs Related To	Details	Configure In
Security events	This syslog configuration generates alerts for file and malware, connection, Security Intelligence, and intrusion events. For details, see *About Sending Syslog Messages for Security Events* and subtopics in the Cisco Secure Firewall Management Center Administration Guide.	**Platform Settings** and the **Logging** in an access control policy
(All devices) Policies, rules, and events	This syslog configuration generates alerts for access control rules, intrusion rules, and other advanced services as described in *Configurations Supporting Alert Responses* in the Cisco Secure Firewall Management Center Administration Guide. These messages are not numbered. For information on configuring this type of syslog, see *Creating a Syslog Alert Response* in the Cisco Secure Firewall Management Center Administration Guide.	**Alert Responses** and the **Logging** in an access control policy

You can configure more than one syslog server, and control the messages and events sent to each server. You can also configure different destinations, such as console, email, internal buffer, and so forth.

Severity Levels

The following table lists the syslog message severity levels.

Table 66: Syslog Message Severity Levels

Level Number	Severity Level	Description
0	**emergencies**	System is unusable.
1	**alert**	Immediate action is needed.
2	**critical**	Critical conditions.
3	**error**	Error conditions.
4	**warning**	Warning conditions.
5	**notification**	Normal but significant conditions.
6	**informational**	Informational messages only.
7	**debugging**	Debugging messages only. Log at this level only temporarily, when debugging issues. This log level can potentially generate so many messages that system performance can be affected.

Note ASA and Firewall Threat Defense do not generate syslog messages with a severity level of zero (emergencies).

Syslog Message Filtering

You can filter generated syslog messages so that only certain syslog messages are sent to a particular output destination. For example, you could configure the Firewall Threat Defense device to send all syslog messages to one output destination and to send a subset of those syslog messages to a different output destination.

Specifically, you can direct syslog messages to an output destination according to the following criteria:

- Syslog message ID number

  (This does not apply to syslog messages for security events such as connection and intrusion events.)

- Syslog message severity level
- Syslog message class (equivalent to a functional area)

  (This does not apply to syslog messages for security events such as connection and intrusion events.)

You customize these criteria by creating a message list that you can specify when you set the output destination. Alternatively, you can configure the Firewall Threat Defense device to send a particular message class to each type of output destination independently of the message list.

(Message lists do not apply to syslog messages for security events such as connection and intrusion events.)

Syslog Message Classes

Note This topic does not apply to messages for security events (connection, intrusion, etc.)

You can use syslog message classes in two ways:

- Specify an output location for an entire category of syslog messages.
- Create a message list that specifies the message class.

The syslog message class provides a method of categorizing syslog messages by type, equivalent to a feature or function of the device. For example, the rip class denotes RIP routing.

All syslog messages in a particular class share the same initial three digits in their syslog message ID numbers. For example, all syslog message IDs that begin with the digits 611 are associated with the vpnc (VPN client) class. Syslog messages associated with the VPN client feature range from 611101 to 611323.

In addition, most of the ISAKMP syslog messages have a common set of prepended objects to help identify the tunnel. These objects precede the descriptive text of a syslog message when available. If the object is not known at the time that the syslog message is generated, the specific heading = value combination does not appear.

The objects are prefixed as follows:

Group = *groupname*, Username = *user*, IP = *IP_address*

Where the group is the tunnel-group, the username is the username from the local database or AAA server, and the IP address is the public IP address of the remote access client or Layer 2 peer.

The following table lists the message classes and the range of message IDs in each class.

Table 67: Syslog Message Classes and Associated Message ID Numbers

Class	Definition	Syslog Message ID Numbers
access-list*	Access Lists	106
application-firewall*	Application Firewall	415
auth	User Authentication	109, 113
botnet-traffic-filtering*	Botnet Traffic Filtering	338
bridge	Transparent Firewall	110, 220
ca	PKI Certification Authority	717
card-management*	Card Management	323
citrix	Citrix Client	723
clustering*	Clustering	747
config	Command Interface	111, 112, 208, 308
csd	Secure Desktop	724
cts	Cisco TrustSec	776
dap	Dynamic Access Policies	734
eap, eapoudp	EAP or EAPoUDP for Network Admission Control	333, 334
eigrp	EIGRP Routing	336
email	E-mail Proxy	719
environment-monitoring*	Environment Monitoring	735
ha	Failover	101, 102, 103, 104, 105, 210, 311, 709
identity-based-firewall*	Identity-based Firewall	746
ids	Intrusion Detection System	400, 733
ikev2-toolkit*	IKEv2 Toolkit	750, 751, 752
ip	IP Stack	209, 215, 313, 317, 408
ipaa	IP Address Assignment	735
ips	Intrusion Protection System	400, 401, 420
ipv6*	IPv6	325
licensing*	Licensing	444

Class	Definition	Syslog Message ID Numbers
mdm-proxy	MDM Proxy	802
nac	Network Admission Control	731, 732
nacpolicy	NAC Policy	731
nacsettings	NAC Settings to apply NAC Policy	732
nat-and-pat*	NAT and PAT	305
network-access-point*	Network Access Point	713
np	Network Processor	319
np-ssl*	NP SSL	725
ospf	OSPF Routing	318, 409, 503, 613
password-encryption*	Password Encryption	742
phone-proxy*	Phone Proxy	337
rip	RIP Routing	107, 312
rm	Resource Manager	321
scansafe*	ScanSafe	775
session	User Session	106, 108, 201, 202, 204, 302, 303, 304, 305, 314, 405, 406, 407, 500, 502, 607, 608, 609, 616, 620, 703, 710
smart-call-home*	Smart Call Home	120
snmp	SNMP	212
ssl	SSL Stack	725
svc	SSL VPN Client	722
sys	System	199, 211, 214, 216, 306, 307, 315, 414, 604, 605, 606, 610, 612, 614, 615,701, 711, 741
tag-switching	Service Tag Switching	779
threat-detection*	Threat Detection	733
transactional-rule-engine-tre*	Transactional Rule Engine	780
uc-ims*	UC-IMS	339
vm	VLAN Mapping	730

Class	Definition	Syslog Message ID Numbers
vpdn	PPTP and L2TP Sessions	213, 403, 603
vpn	IKE and IPsec	316, 320, 402, 404, 501, 602, 702, 713, 714, 715
vpnc	VPN Client	611
vpnfo	VPN Failover	720
vpnlb	VPN Load Balancing	718
vxlan*	VXLAN	778
webfo	WebVPN Failover	721
webvpn	WebVPN and Secure Client	716

*These classes are provisioned in the management center web interface to facilitate creation of event lists. These classes are not displayed on the device console.

Guidelines for Logging

This section includes guidelines and limitations that you should review before configuring logging.

IPv6 Guidelines

- IPv6 is supported. Syslogs can be sent using TCP or UDP.
- Ensure that the interface configured for sending syslogs is enabled, IPv6 capable, and the syslog server is reachable through the designated interface.
- Secure logging over IPv6 is not supported.

Additional Guidelines

- Do not configure Firewall Management Center as a primary syslog server. The Firewall Management Center can log some syslogs. However, it does not have adequate storage provision to accommodate voluminous information from connection events for every sensor, especially when multiple sensors are used and all send syslogs.
- The syslog server must run a server program called syslogd. Windows provides a syslog server as part of its operating system.
- The syslog server operates based on the syslog-ng process of the firewall system. Do not use external configuration files, like the *scwx.conf* file from SecureWorks. Such files are not compatible with the device. Using them will lead to parsing error and eventually the syslog-ng process will fail.
- To view logs generated by the Firewall Threat Defense device, you must specify a logging output destination. If you enable logging without specifying a logging output destination, the Firewall Threat Defense device generates messages but does not save them to a location from which you can view them. You must specify each different logging output destination separately.

- If you use TCP as the transport protocol, the system opens 4 connections to the syslog server to ensure that messages are not lost. If you are using the syslog server to collect messages from a very large number of devices, and the combined connection overhead is too much for the server, use UDP instead.
- It is not possible to have two different lists or classes being assigned to different syslog servers or same locations.
- You can configure up to 16 syslog servers.
- The syslog server should be reachable through the Firewall Threat Defense device. You should configure the device to deny ICMP unreachable messages on the interface through which the syslog server is reachable and to send syslogs to the same server. Make sure that you have enabled logging for all severity levels. To prevent the syslog server from crashing, suppress the generation of syslogs 313001, 313004, and 313005.
- The number of UDP connections for syslog is directly related to the number of CPUs on the hardware platform and the number of syslog servers you configure. At any point in time, there can be as many UDP syslog connections as there are CPUs times the number of configured syslog servers. This is the expected behavior. Note that the global UDP connection idle timeout applies to these sessions, and the default is 2 minutes. You can adjust that setting if you want to close these session more quickly, but the timeout applies to all UDP connections, not just syslog.
- When the Firewall Threat Defense device sends syslogs via TCP, the connection takes about one minute to initiate after the syslogd service restarts.
- When the TCP logging host goes down, it takes approximately 6 minutes to change its connection status from *Connected* to *Not connected*. Logging relies on TCP to detect the channel state; until then, logging sends the logs through the channel. During this time, when you execute the **show log**, the output would display the TCP logging host as connected. Once the TCP channel is closed, the TCP logging host state is updated to *Not connected*.

Configure Syslog Logging for Firewall Threat Defense Devices

Tip If you are configuring devices to send syslog messages about security events (such as connection and intrusion events), most Firewall Threat Defense platform settings do not apply to these messages. See Firewall Threat Defense Platform Settings That Apply to Security Event Syslog Messages, on page 991.

To configure syslog settings, perform the following steps:

Before you begin

See requirements in Guidelines for Logging, on page 989.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Click **Syslog** from the table of contents.

Step 3 Click **Logging Setup** to enable logging, specify FTP Server settings, and specify Flash usage. For more information, see Enable Logging and Configure Basic Settings, on page 991

Step 4 Click **Logging Destinations** to enable logging to specific destinations and to specify filtering on message severity level, event class, or on a custom event list. For more information, see Enable Logging Destinations, on page 993

You must enable a logging destination to see messages at that destination.

Step 5 Click **E-mail Setup** to specify the e-mail address that is used as the source address for syslog messages that are sent as e-mail messages. For more information, see Send Syslog Messages to an E-mail Address, on page 994

Step 6 Click **Events List** to define a custom event list that includes an event class, a severity level, and an event ID. For more information, see Create a Custom Event List, on page 995

Step 7 Click **Rate Limit** to specify the volume of messages being sent to all configured destinations and define the message severity level to which you want to assign rate limits. For more information, see Limit the Rate of Syslog Message Generation, on page 996

Step 8 Click **Syslog Settings** to specify the logging facility, enable the inclusion of a time stamp, and enable other settings to set up a server as a syslog destination. For more information, see Configure Syslog Settings, on page 997

Step 9 Click **Syslog Servers** to specify the IP address, protocol used, format, and security zone for the syslog server that is designated as a logging destination. For more information, see Configure a Syslog Server, on page 998.

Firewall Threat Defense Platform Settings That Apply to Security Event Syslog Messages

"Security events" include connection, Security Intelligence, intrusion, and file and malware events.

Some of the syslog settings configured on the **Syslog** page (**Devices** > **Platform Settings**, edit Threat Defense Policy, and then click **Syslog**) and its tabs apply to syslog messages for security events, but most apply only to messages for events related to system health and networking.

The following settings apply to syslog messages for security events:

- **Logging Setup** tab:
  - **Send syslogs in EMBLEM format**
- **Syslog Settings** tab:
  - **Enable Timestamp on Syslog Messages**
  - **Timestamp Format**
  - **Enable Syslog Device ID**
- **Syslog Servers** tab:
  - All options on the **Add Syslog Server** form (and the list of configured servers).

Enable Logging and Configure Basic Settings

Enable logging and configure the basic settings for the system to generate syslog messages for data plane events. You can also set up archiving on flash or an FTP server as a storage location when the local buffer

becomes full. You can manipulate logging data after it is saved. For example, you could specify actions to be executed when certain types of syslog messages are logged, extract data from the log and save the records to another file for reporting, or track statistics using a site-specific script.

The following procedure explains some of the basic syslog settings.

Tip If you are configuring devices to send syslog messages about security events (such as connection and intrusion events), most Firewall Threat Defense platform settings do not apply to these messages. See Firewall Threat Defense Platform Settings That Apply to Security Event Syslog Messages, on page 991.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Syslog** > **Logging Setup**.

Step 3 Enable logging and configure basic logging settings.

- **Enable Logging**—Turns on the data plane system logging for the Firewall Threat Defense device.
- **Enable Logging on the Failover Standby Unit**—Turns on logging for the standby for the Firewall Threat Defense device, if available.
- **Send syslogs in EMBLEM format**—Enables EMBLEM format logging for every logging destination. If you enable EMBLEM, you must use the UDP protocol to publish syslog messages; EMBLEM is not compatible with TCP.

  The EMBLEM syslog format is a Cisco-specific convention that is built upon the RFC 3164 and RFC 5424 standards. Hence, when EMBLEM is enabled, the syslog message prints colon (:) after <PRI> field in contrast to the RFC 5424 format.

  Note
  Syslog messages in RFC5424 format, typically displays the priority value (PRI). However, in Firewall Management Center, if you want to display the PRI value in the syslog messages of the managed Firewall Threat Defense device, ensure to enable the EMBLEM format. Also, the device ID does not appear in EMBLEM-formatted syslog messages.

- **Send debug messages as syslogs**—Redirects all the debug trace output to the syslog. The syslog message does not appear in the console if this option is enabled. Therefore, to see debug messages, you must enable logging at the console and configure it as the destination for the debug syslog message number and logging level. The syslog message number used is 711001. The default logging level for this syslog is debug.
- **Memory Size of Internal Buffer**—Specify the size of the internal buffer to which syslog messages are saved if the logging buffer is enabled. When the buffer fills up, it is overwritten. The default is 4096 bytes. The range is 4096 to 52428800.

Step 4 (Optional) Configure the syslog message logging to the Firewall Management Center.

a) Click the **All Logs** radio button to enable logging all the troubleshooting syslog messages corresponding to the selected severity level or click the **VPN Logs** radio button to enable logging only the VPN troubleshooting messages corresponding to the selected severity level.

b) Choose the syslog severity level for the logging messages from the **Logging Level** drop-down list.

- The logging level for **All Logs** is set to **critical** by default. You can choose to send syslog messages with severity levels **critical**, **alerts**, or **emergencies** to the Firewall Management Center.
- The logging level for the VPN messages is set to **errors** by default.

  VPN troubleshooting syslogs can add excessive load on the Firewall Management Center. Hence, enable this option with caution. Also, when you configure a device with site-to-site or remote access VPN, it automatically enables sending VPN syslogs to the management center by default. We recommend that you limit the logging level to **error** and above to restrict the excessive flow of syslogs to the Firewall Management Center, especially in case of RAVPN, where multiple devices are involved.

For information on the levels, see Severity Levels, on page 985.

Step 5 (Optional) Configure an FTP server if you want to save log buffer contents to the server before the buffer is overwritten. Specify the FTP Server information.

- **FTP Server Buffer Wrap**—To save the buffer contents to the FTP server before it is overwritten, check this box and enter the necessary destination information in the following fields. To remove the FTP configuration, deselect this option.
- **IP Address**—Select the host network object that contains the IP address of the FTP server.
- **User Name**—Enter the username to use when connecting to the FTP server.
- **Path**—Enter the path, relative to the FTP root, where the buffer contents should be saved.
- **Password/ Confirm**—Enter and confirm the password used to authenticate the username to the FTP server.

Step 6 (Optional) Specify Flash size if you want to save log buffer contents to flash before the buffer is overwritten.

- **Flash**—To save the buffer contents to the flash memory before it is overwritten, check this box.
- **Maximum flash to be used by logging (KB)**—Specify the maximum space to be used in the flash memory for logging (in kilobytes). The range is 4-8044176 kilobytes.
- **Minimum free space to be preserved (KB)**—Specifies the minimum free space to be preserved in flash memory (in KB). The range is 0-8044176 kilobytes.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Enable Logging Destinations

You must enable a logging destination to see messages at that destination. When enabling a destination, you must also specify the message filter for the destination.

Tip If you are configuring devices to send syslog messages about security events (such as connection and intrusion events), most Firewall Threat Defense platform settings do not apply to these messages. See Firewall Threat Defense Platform Settings That Apply to Security Event Syslog Messages, on page 991.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Syslog** > **Logging Destinations**.

Step 3 Click **Add** to enable a destination and apply a logging filter, or edit an existing destination.

Step 4 In the **Logging Destinations** dialog box, select a destination and configure the filter to use for a destination:

a) Choose the destination you are enabling in the **Logging Destination** drop-down list. You can create one filter per destination: Console, E-Mail, Internal buffer, SNMP trap, SSH Sessions, and Syslog servers.

Note
Console and SSH session logging works in the diagnostic CLI only. Enter **system support diagnostic-cli**.

b) In **Event Class**, choose the filter that will apply to all classes not listed in the table.

You can configure these filters:

- **Filter on severity** —Select the severity level. Messages at this level or higher are sent to the destination
- **Use Event List** —Select the event list that defines the filter. You create these lists on the **Event Lists** page.
- **Disable Logging** —Prevents messages from being sent to this destination.

c) If you want to create filters per event class, click **Add** to create a new filter, or edit an existing filter, and select the event class and severity level to limit messages in that class. Click **OK** to save the filter.

For an explanation of the event classes, see Syslog Message Classes, on page 986.

d) Click **OK** .

Step 5 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Send Syslog Messages to an E-mail Address

You can set up a list of recipients for syslog messages to be sent as e-mails.

Tip If you are configuring devices to send syslog messages about security events (such as connection and intrusion events), most Firewall Threat Defense platform settings do not apply to these messages. See Firewall Threat Defense Platform Settings That Apply to Security Event Syslog Messages, on page 991.

Before you begin

- Configure an SMTP server on the SMTP Server platform settings page
- Enable Logging and Configure Basic Settings, on page 991

- Setup Email as the Logging Destination

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Syslog** > **Email Setup**.

Step 3 Specify the e-mail address that is used as the source address for syslog messages that are sent as e-mail messages.

Step 4 Click **Add** to enter a new e-mail address recipient of the specified syslog messages.

Step 5 Choose the severity level of the syslog messages that are sent to the recipient from the drop-down list.

The syslog message severity filter used for the destination e-mail address causes messages of the specified severity level and higher to be sent. For information on the levels, see Severity Levels, on page 985.

Step 6 Click **OK**.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Create a Custom Event List

An event list is a custom filter you can apply to a logging destination to control which messages are sent to the destination. Normally, you filter messages for a destination based on severity only, but you can use an event list to fine-tune which messages are sent based on a combination of event class, severity, and message identifier (ID).

Creating a custom event list is a two-step process. You create a custom list in the **Event Lists**, and then use the event list to define the logging filter for the various types of destination, in the **Logging Destinations**.

Tip If you are configuring devices to send syslog messages about security events (such as connection and intrusion events), most Firewall Threat Defense platform settings do not apply to these messages. See Firewall Threat Defense Platform Settings That Apply to Security Event Syslog Messages, on page 991.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Syslog** > **Events List**.

Step 3 Configure an event list.

a) Click **Add** to add a new list, or edit an existing list.

b) Enter a name for the event list in the **Name** field. Spaces are not allowed.

c) To identify messages based on severity or event class, select the **Severity/Event Class** tab and add or edit entries.

For information on the available classes see Syslog Message Classes, on page 986.

For information on the levels, see Severity Levels, on page 985.

Certain event classes are not applicable for the device in transparent mode. If such options are configured then they will be bypassed and not deployed.

d) To identify messages specifically by message ID, select the **Message ID** and add or edit the IDs.

You can enter a range of IDs using a hyphen, for example, 100000-200000. IDs are six digits. For information on how the initial three digits map to features, see Syslog Message Classes, on page 986.

For specific message numbers, see Cisco ASA Series Syslog Messages.

e) Click **OK** to save the event list.

Step 4 Click **Logging Destinations** and add or edit the destination that should use the filter.

See Enable Logging Destinations, on page 993.

Step 5 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Limit the Rate of Syslog Message Generation

You can limit the rate at which syslog messages are generated by severity level or message ID. You can specify individual limits for each logging level and each Syslog message ID. If the settings conflict, the Syslog message ID limits take precedence.

Tip If you are configuring devices to send syslog messages about security events (such as connection and intrusion events), most Firewall Threat Defense platform settings do not apply to these messages. See Firewall Threat Defense Platform Settings That Apply to Security Event Syslog Messages, on page 991.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Syslog** > **Rate Limit**.

Step 3 To limit message generation by severity level, click **Logging Level** > **Add** and configure the following options:

- **Logging Level**—The severity level you are rate limiting. For information on the levels, see Severity Levels, on page 985.
- **Number of messages**—The maximum number of messages of the specified type allowed in the specified time period.
- **Interval**—The number of seconds before the rate limit counter resets.

Step 4 Click **OK**.

Step 5 To limit message generation by syslog message ID, click **Syslog Level** > **Add** and configure the following options:

- **Syslog ID**—The syslog message ID you are rate limiting. For specific message numbers, see Cisco ASA Series Syslog Messages.
- **Number of messages**—The maximum number of messages of the specified type allowed in the specified time period.
- **Interval**—The number of seconds before the rate limit counter resets.

Step 6 Click **OK**.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Syslog Settings

You can configure general syslog settings to set the facility code to be included in syslog messages that are sent to syslog servers, specify whether a timestamp is included in each message, specify the device ID to include in messages, view and modify the severity levels for messages, and disable the generation of specific messages.

If you are configuring devices to send syslog messages about security events (such as connection and intrusion events), some settings on this page do not apply to these messages. See *Threat Defense Platform Settings That Apply to Security Event Syslog Messages* in the Cisco Secure Firewall Management Center Administration Guide.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Syslog** > **Syslog Settings**.

Step 3 Select a system log facility for syslog servers to use as a basis to file messages in the **Facility** drop-down list.

The default is LOCAL4(20), which is what most UNIX systems expect. However, because your network devices share available facilities, you might need to change this value for system logs.

Facility values are not typically relevant for security events. If you need to include Facility values in messages, see *Facility in Security Event Syslog Messages* in the Cisco Secure Firewall Management Center Administration Guide.

Step 4 Select the **Enable timestamp on each syslog message** check box to include the date and time a message was generated in the syslog message.

Step 5 Select the **Timestamp Format** for the syslog message:

- The Legacy (MMM dd yyyy HH:mm:ss) format is the default format for syslog messages.

  When this timestamp format is selected, the messages do not indicate the time zone, which is always UTC.

- RFC 5424 (yyyy-MM-ddTHH:mm:ssZ) uses the ISO 8601 timestamp format as specified in the RFC 5424 syslog format.

  If you select the RFC 5424 format, a “Z” is appended to the end of each timestamp to indicate that the timestamp uses the UTC time zone.

Step 6 If you want to add a device identifier to syslog messages (which is placed at the beginning of the message), check the **Enable Syslog Device ID** check box and then select the type of ID.

- **Interface**—To use the IP address of the selected interface, regardless of the interface through which the appliance sends the message. Select the security zone that identifies the interface. The zone must map to a single interface.
- **User Defined ID**—To use a text string (up to 16 characters) of your choice.
- **Host Name**—To use the hostname of the device.

Step 7 Use the Syslog Message table to alter the default settings for specific syslog messages. You need to configure rules in this table only if you want to change the default settings. You can change the severity assigned to a message, or you can disable the generation of a message.

By default, Netflow is enabled and the entries are shown in the table.

a) To suppress syslog messages that are redundant because of Netflow, select **Netflow Equivalent Syslogs**.

This adds the messages to the table as suppressed messages.

Note
If any of these syslog equivalents are already in the table, your existing rules are not overwritten.

b) To add a rule, click **Add**.
c) You select the message number whose configuration you want to change, from the **Syslog ID** drop down list and then select the new severity level from the **Logging Level** drop down list, or select **Suppressed** to disable the generation of the message. Typically, you would not change the severity level and disable the message, but you can make changes to both fields if desired.
d) Click **OK** to add the rule to the table.

Step 8 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Configure a Syslog Server

To configure a syslog server to handle messages generated from your system, perform the following steps.

If you want this syslog server to receive security events such as connection and intrusion events, see also Firewall Threat Defense Platform Settings That Apply to Security Event Syslog Messages, on page 991.

Note In 7.4 and later, the Management and Diagnostic interfaces are merged. If Platform Settings for syslog servers or SNMP hosts specify the Diagnostic interface by name, then you must use separate Platform Settings policies for merged and unmerged devices (7.3 and earlier, and some upgraded 7.4 Firewall Threat Defenses).

Before you begin

- See requirements in Guidelines for Logging, on page 989.
- Make sure your devices can reach your syslog collector on the network.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Syslog** > **Syslog Server**.

Step 3 Check the **Allow user traffic to pass when TCP syslog server is down (Recommended)** check box, to allow traffic if any syslog server that is using the TCP protocol is down.

Note

- This option is enabled by default. Unless required, we recommend that you allow connections through the threat defense device when the external TCP syslog server is unreachable by the device.
- When the **Allow user traffic to pass when TCP syslog server is down** option is disabled in Firewall Management Center version 6.2.x or earlier, it persist to be in the Disable state even after upgrading to version 6.3 or later. Ensure that you manually enable it.
- With this option disabled, and when more than one TCP syslog server configured in the device, the user traffic is allowed to pass if atleast one of the servers is reacheable by the threat defense device. Thus, the disabled option is applied only when none of the TCP syslog servers configured in the device are reachable. The device generates the following syslog that describes the root cause of the denied traffic through the device:

```
%FTD-3-414003: TCP Syslog Server intf : IP_Address /port not responding. New connections
are denied based on logging permit-hostdown policy
```

Step 4 In the **Message queue size (messages)** field, enter a size of the queue for storing syslog messages on the security appliance when syslog server is busy. The minimum is 1 message. The default is 512. Specify 0 to allow an unlimited number of messages to be queued (subject to available block memory).

When the messages exceed the configured queue size, they are dropped and result in missing syslog. To determine the ideal queue size, you need to identify the available block memory. Use the **show blocks** command to know the current memory utilization. For more information on the command and its attributes, see *Cisco Secure Firewall ASA Series Command Reference Guide*. For further assistance, contact Cisco TAC.

Step 5 Click **Add** to add a new syslog server.

a) In the **IP Address** drop-down list, select a network host object that contains the IP address of the syslog server.

b) Choose the protocol (either TCP or UDP) and enter the port number for communications between the Firewall Threat Defense device and the syslog server.

UDP is faster and uses less resources on the device than TCP.

The default port for UDP is 514. You must manually configure port 1470 for TCP. Valid non-default port values for either protocol are 1025 through 65535.

c) Check the **Log messages in Cisco EMBLEM format (UDP only)** check box to specify whether to log messages in Cisco EMBLEM format (available only if UDP is selected as the protocol).

The EMBLEM syslog format is a Cisco-specific convention that is built upon the RFC 3164 and RFC 5424 standards. Hence, when EMBLEM is enabled, the syslog message prints colon (:) after <PRI> field.

Note
Syslog messages in RFC5424 format, typically displays the priority value (PRI). However, in Firewall Management Center, only when you enable logging in Cisco EMBLEM format, the PRI value in the syslog messages of the managed Firewall Threat Defense is displayed. Also, the device ID does not appear in EMBLEM-formatted syslog messages.

d) Check the **Enable Secure Syslog** check box to encrypt the connection between the device and server using SSL/TLS over TCP.

Note
You must select TCP as the protocol and its port value ranging between 1025 and 65535 to use this option. You must also upload the certificate required to communicate with the syslog server on the **Devices** > **Certificates** page. Finally, upload the certificate from the Firewall Threat Defense device to the syslog server to complete the secure relationship and allow it to decrypt the traffic. The **Enable Secure Syslog** option is not supported on the device Management interface.

e) Select **Device Management Interface** or **Security Zones or Named Interfaces** to communicate with the syslog server.

- **Device Management Interface**: Send syslogs out of the Management interface. We recommend that you use this option when configuring syslog on Snort events.

  Note
  The **Device Management Interface** option does not support the **Enable Secure Syslog** option.

- **Security Zones or Named Interfaces**: Select the interfaces from the list of **Available Zones** and click **Add**. You can also add virtual-router-aware interfaces.

  Important
  The Firewall Threat Defense data plane (Lina) syslog messages cannot be sent out through the diagnostic interface. Configure other interfaces or the Management interface (Br1/Management0) to send out the data plane syslog messages.

f) Click **OK**.

Step 6 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Timeouts

You can set the global idle timeout durations for the connection and translation slots of various protocols. If the slot has not been used for the idle time specified, the resource is returned to the free pool.

You can also set a time out for console sessions with the device.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Timeouts**.

Step 3 Configure the timeouts you want to change.

For any given setting, select **Custom** to define your own value, **Default** to return to the system default value. In most cases, the maximum timeout is 1193 hours.

You can disable some timeouts by selecting **Disable**.

- **Console Timeout**—The idle time until a connection to the console is closed, range is 0 or 5 to 1440 minutes. The default is 0, which means the session does not time out. If you change the value, existing console sessions use the old timeout value. The new value applies to new connections only.
- **Translation Slot (xlate)**—The idle time until a NAT translation slot is freed. This duration must be at least 1 minute. The default is 3 hours.
- **Connection (Conn)**—The idle time until a connection slot is freed. This duration must be at least 5 minutes. The default is 1 hour.
- **Half-Closed**—The idle time until a TCP half-closed connection closes. A connection is considered half-closed if both the FIN and FIN-ACK have been seen. If only the FIN has been seen, the regular connection timeout applies. The minimum is 30 seconds. The default is 10 minutes.
- **UDP**—The idle time until a UDP connection closes. This duration must be at least 1 minute. The default is 2 minutes.
- **ICMP**—The idle time after which general ICMP states are closed. The default (and minimum) is 2 seconds.
- **RPC/Sun RPC**—The idle time until a SunRPC slot is freed. This duration must be at least 1 minute. The default is 10 minutes.

  In a Sun RPC-based connection, when the parent connection is deleted or timed-out, a new child connection may not be considered as a part of the parent-child connection, and thereby the new connection could be evaluated as per the policy or rules set in the system. After the parent connection has timed-out the existing child connections are valid only until the timeout value set is reached.
- **H.225**—The idle time until an H.225 signaling connection closes. The default is 1 hour. To close a connection immediately after all calls are cleared, a timeout of 1 second (0:0:1) is recommended.
- **H.323**—The idle time after which H.245 (TCP) and H.323 (UDP) media connections close. The default (and minimum) is 5 minutes. Because the same connection flag is set on both H.245 and H.323 media connections, the H.245 (TCP) connection shares the idle timeout with the H.323 (RTP and RTCP) media connection.
- **SIP**—The idle time until a SIP signaling port connection closes. This duration must be at least 5 minutes. The default is 30 minutes.
- **SIP Media**—The idle time until a SIP media port connection closes. This duration must be at least 1 minute. The default is 2 minutes. The SIP media timer is used for SIP RTP/RTCP with SIP UDP media packets, instead of the UDP inactivity timeout.

- **SIP Disconnect**—The idle time after which SIP session is deleted if the 200 OK is not received for a CANCEL or a BYE message, between 0:0:1 and 0:10:0. The default is 2 minutes (0:2:0).
- **SIP Invite**—The idle time after which pinholes for PROVISIONAL responses and media xlates will be closed, between 0:1:0 and 00:30:0. The default is 3 minutes (0:3:0).
- **SIP Provisional Media**—The timeout value for SIP provisional media connections, between 1 and 30 minutes. The default is 2 minutes.
- **Floating Connection**—When multiple routes exist to a network with different metrics, the system uses the one with the best metric at the time of connection creation. If a better route becomes available, then this timeout lets connections be closed so a connection can be reestablished to use the better route. The default is 0 (the connection never times out). To make it possible to use better routes, set the timeout to a value between 0:0:30 and 1193:0:0. This timer does not apply to connections through virtual tunnel interfaces (VTI). If a connection through a VTI gets stuck, you must manually clear it.
- **Xlate PAT**—The idle time until a PAT translation slot is freed, between 0:0:30 and 0:5:0. The default is 30 seconds. You may want to increase the timeout if upstream routers reject new connections using a freed PAT port because the previous connection might still be open on the upstream device.
- **TCP Proxy Reassembly**—The idle timeout after which buffered packets waiting for reassembly are dropped, between 0:0:10 and 1193:0:0. The default is 1 minute (0:1:0).
- **ARP Timeout**—The number of seconds between ARP table rebuilds, from 60 to 4294967. The default is 14,400 seconds (4 hours).

Step 4 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Time Synchronization

Use a Network Time Protocol (NTP) server to synchronize the clock settings on your devices. We recommend you configure all Firewall Threat Defenses managed by a Firewall Management Center to use the same NTP server as the Firewall Management Center. The Firewall Threat Defense gets its time directly from the configured NTP server. If the Firewall Threat Defense's configured NTP servers are not reachable for any reason, it synchronizes its time with the Firewall Management Center.

The device supports NTPv4.

Note If you are deploying Firewall Threat Defense on the Firepower 4100/9300 chassis, you must configure NTP on the Firepower 4100/9300 chassis so that Smart Licensing will work properly and to ensure proper timestamps on device registrations. You should use the same NTP server for the Firepower 4100/9300 chassis and the Firewall Management Center.

Before you begin

- If your organization has one or more NTP servers that your Firewall Threat Defense can reach, use the same NTP server or servers for your devices that you have configured for Time Synchronization on the **System** (⚙) > **Configuration** page on your Firewall Management Center.
- If you selected **Use the authenticated NTP server only** when configuring NTP server or servers for the Firewall Management Center, for your devices use only the NTP server or servers that are configured to authenticate with the Firewall Management Center. (The managed devices will use the same NTP servers as the Firewall Management Center, but their NTP connections will not use authentication.)
- If your device cannot reach an NTP server or your organization does not have one, you must use the **Via NTP from Defense Center** option as discussed in the following procedure.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Time Synchronization**.

Step 3 Configure one of the following clock options:

- **Via NTP from Defense Center**—(Default). The managed device gets time from the NTP servers you configured for the Firewall Management Center (except for authenticated NTP servers) and synchronizes time with those servers directly. However, if any of the following are true, the managed device synchronizes time from the Firewall Management Center:
  - The Firewall Management Center's NTP servers are not reachable by the device.
  - The Firewall Management Center has no unauthenticated servers.
- **Via NTP from**—If your Firewall Management Center is using NTP servers on the network, select this option and enter the fully-qualified DNS name (such as ntp.example.com), or IPv4 or IPv6 address, of the same NTP servers you specified in **System** (⚙) > **Configuration** > **Time Synchronization**. If the NTP servers are not reachable, the Firewall Management Center acts as an NTP server.

When multiple NTP servers are configured, the device uses the NTP server that is deemed appropriate based on the criteria defined in RFC. Thus, the status of "Being used" for a specific NTP server indicates that the server is currently used by the device.

Step 4 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Time Zone

By default, the system uses the UTC time zone. To designate a different time zone for a device, use this procedure.

The time zone you specify will be used only for time-based policy application in policies that support this functionality.

Note Time-based ACLs is supported in Snort 3 also from Firewall Management Center 7.0 onwards.

Procedure

Step 1 Select **Devices** > **Platform Settings** and create or edit an Firewall Threat Defense policy.

You can also create time zone objects from the **Objects** > **Object Management** > **Time Zone** page.

Step 2 Create a new time zone object by clicking **+**.

Step 3 Select the time zone.

Step 4 Click **Save**.

What to do next

- Create time range objects, select applicable time ranges in access control and prefilter rules, and assign the parent policies to devices associated with the correct time zone.
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

UCAPL/CC Compliance

For more information about this setting and how to enable it for the Firewall Management Center, see the Cisco Secure Firewall Management Center Administration Guide.

Caution After you enable this setting, you cannot disable it. If you need to take the appliance out of CC or UCAPL mode, you must reimage.

Before you begin

- Secure Firewall Threat Defense devices cannot use an evaluation license; your Smart Software Manager account must be enabled for export-controlled features.
- Secure Firewall Threat Defense devices must be deployed in routed mode.
- You must be an Admin user to perform this task.
- Security certifications compliance mode cannot be changed if the Firewall Threat Defense device is in high availability. Modify the security certifications compliance mode before forming the high availability pair.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Click **UCAPL/CC Compliance**.

Step 3 To *permanently* enable security certifications compliance on the appliance, you have two choices:

- To enable security certifications compliance in Common Criteria mode, choose **CC** from the drop-down list.
- To enable security certifications compliance in Unified Capabilities Approved Products List mode, choose **UCAPL** from the drop-down list.

Step 4 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Performance Profile

The performance profile determines how the CPU cores on the device are assigned to two of the main system processes: the data plane (Lina) and Snort. The data plane handles VPN connections, routing, and other basic layer 3/4 processing. Snort provides advanced inspection, including intrusion and malware prevention, URL filtering, application filtering, and other features that require deep packet inspection.

If you use a balance of basic and advanced features, do not change the performance profile. The system is designed to provide a balanced assignment of cores to these processes. The assignment differs based on the hardware model.

However, if you use the device primarily for VPN, or for intrusion and other advanced inspection, you can skew the performance profile so that more cores are assigned to the more heavily used features. This might improve system performance.

Before you begin

- These settings apply to systems running release 7.3+ only.
- Performance profile is supported on the following device types:
  - Firepower 4100/9300
  - Secure Firewall 3100/4200 (7.4+)
  - Secure Firewall Threat Defense Virtual
- Changing the performance profile is not supported on units in a cluster or high-availability group, or those configured for multi-instance. Deployment is blocked if you assign the profile to anything but standalone devices.
- The minimum number for core allocation is 2. Cores are assigned in even numbers based on the selected performance profile.

Procedure

Step 1 Choose **Devices** > **Platform Settings** and create or edit the Firewall Threat Defense policy.

Step 2 Select **Performance Profile**.

Step 3 Select a profile:

- **Default**—This is the recommended setting and is the best option if you configure both VPN and intrusion inspection.
- **VPN Heavy with prefilter fastpath**—If you primarily use the device as a VPN endpoint or headend, and you configure rules in the prefilter policy to fastpath VPN traffic, you can choose this option to assign the majority of CPU cores to the data plane. The allocation is 90% data place, 10% Snort.
- **VPN Heavy with inspection**—If you primarily use the device as a VPN endpoint or headend, but do not use the prefilter policy to fastpath VPN traffic, you can choose this option to assign the majority of CPU cores to the data plane. This option assumes that you leave intrusion inspection, URL filtering, and other advanced functions that use Snort, to a different device in the network. The allocation is 60% data plane, 40% Snort.
- **IPS Heavy**—If you do not configure VPN, but you do use the device for intrusion prevention, you can choose this option to assign the majority of CPU core to the Snort process. The allocation is 30% data plane, 70% Snort.

Step 4 Click **Save**.

Step 5 Deploy the policy.

Step 6 After deployment completes, you must reboot each affected device so that the new core assignments can be made.

History for Platform Settings

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Device management services supported on user-defined VRF interfaces.	7.6.0	7.6.0	Device access management is now supported on user-defined Virtual Routing and Forwarding (VRF) interfaces. See: Platform Settings
View diagnostic syslog messages in the Unified Events table.	7.6.0	Any	You can now configure Firewall Threat Defense devices to send all device syslogs to your Firewall Management Center. Go to the Firewall Threat Defense platform settings policy assigned to your managed devices and select **All Logs** under the **Logging to Secure Firewall Management Center** setting. Once you have enabled the logging settings, you can view the device syslogs as a new event type called **Troubleshoot Events** in the **Unified Events** table.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Support for Firepower Platform Settings policy is removed.	7.6.0	7.6.0	The provision to manage a Firepower Platform Settings policy is removed.
Device management services supported on user-defined VRF interfaces.	7.4.1	7.4.1	Device management services configured in the Firewall Threat Defense platform settings (NetFlow, SSH access, SNMP hosts, syslog servers) are now supported on user-defined Virtual Routing and Forwarding (VRF) interfaces. Platform restrictions: Not supported with container instances or clustered devices.
Chassis platform settings for the Secure Firewall 3100.	7.4.1	7.4.1	New platform settings for multi-instance chassis on the Secure Firewall 3100. New/modified screens: • **Devices** > **Platform Settings** > **Add** > **Chassis Platform Settings** • **Devices** > **Platform Settings** > **Edit** > **Chassis Platform Settings** > **DNS** • **Devices** > **Platform Settings** > **Edit** > **Chassis Platform Settings** > **SSH** • **Devices** > **Platform Settings** > **Edit** > **Chassis Platform Settings** > **Time Synchronization** • **Devices** > **Platform Settings** > **Edit** > **Chassis Platform Settings** > **Time Zones** • **Devices** > **Platform Settings** > **Edit** > **Chassis Platform Settings** > **Syslog** Supported platforms: Secure Firewall 3100
Performance profile support for the Secure Firewall 3100/4200.	7.4.0	7.4.0	The performance profile settings available in the platform settings policy now apply to Secure Firewall 3100/4200 devices.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Loopback interface support for DNS, HTTP, ICMP, NetFlow, SNMP and SSH.	7.4.0	7.4.0	You can create a loopback interface and use it for: • DNS • HTTP • ICMP • NetFlow • SNMP • SSH New/modified screens: • **Devices** > **Platform Settings** > **DNS** > **DNS Settings** **Devices** > **Platform Settings** > **HTTP Access** > **Add** **Devices** > **Platform Settings** > **ICMP Access** > **Add** • **Devices** > **Platform Settings** > **Network flow** > **Add Collector** • **Devices** > **Platform Settings** > **SNMP** > **Hosts** > **Add** • **Devices** > **Platform Settings** > **SSH Access** > **Add**

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Merged Management and Diagnostic interfaces.	7.4.0	7.4.0	For new devices using 7.4 and later, you cannot use the legacy Diagnostic interface. Only the merged Management interface is available. If you upgraded to 7.4 or later, and you did not have any configuration for the Diagnostic interface, then the interfaces will merge automatically. If you upgraded to 7.4 or later, and you have configuration for the Diagnostic interface, then you have the choice to merge the interfaces manually, or you can continue to use the separate Diagnostic interface. Note that support for the Diagnostic interface will be removed in a later release, so you should plan to merge the interfaces as soon as possible. Merged mode also changes the behavior of AAA traffic to use the data routing table by default. The management-only routing table can now only be used if you specify the management-only interface (including Management) in the configuration. For platform settings, this means: • You can no longer enable HTTP, ICMP, or SMTP for Diagnostic. • For SNMP, you can allow hosts on Management instead of Diagnostic. • For Syslog servers, you can reach them on Management instead of Diagnostic. • If Platform Settings for syslog servers or SNMP hosts specify the Diagnostic interface by name, then you must use separate Platform Settings policies for merged and non-merged devices. • DNS lookups no longer fall back to the management-only routing table if you do not specify interfaces. New/modified screens: • **Devices** > **Device Management** > **Interfaces** New/modified commands: **show management-interface convergence**
Performance profile for CPU core allocation.	7.3.0	7.3.0	You can adjust the percentage of system cores assigned to the data plane and Snort to adjust system performance. The adjustment is based on your relative use of VPN and intrusion policies. If you use both, leave the core allocation to the default values. If you use the system primarily for VPN (without applying intrusion policies), or as an IPS (with no VPN configuration), you can skew the core allocation to the data plane (for VPN) or Snort (for intrusion inspection). We added the **Performance Profile** page to the platform settings policy.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
TLS 1.3 in Remote Access VPN.	7.3.0	7.3.0	You can now use TLS 1.3 to encrypt remote access VPN connections. Use threat defense platform settings to specify that the device must use TLS 1.3 protocol when acting as a remote access VPN server. TLS 1.3 adds support for the following ciphers: • TLS_AES_128_GCM_SHA256 • TLS_CHACHA20_POLY1305_SHA256 • TLS_AES_256_GCM_SHA384 This feature requires Cisco Secure Client, Version 5.0 and above. New/modified screens: **Devices** > **Platform Settings** > **New Policy** > **Add/Edit Threat Defense Settings** > **SSL** > **TLS version**
Multiple DNS server groups for resolving DNS requests.	7.2.0	7.2.0	You can configure multiple DNS groups for the resolution of DNS requests from client systems. You can use these DNS server groups to resolve requests for different DNS domains. For example, you could have a catch-all default group that uses public DNS servers, for use with connections to the Internet. You could then configure a separate group to use internal DNS servers for internal traffic, for example, any connection to a machine in the example.com domain. Thus, connections to an FQDN using your organization's domain name would be resolved using your internal DNS servers, whereas connections to public servers use external DNS servers. We changed the **Platform Settings** > **DNS** page.
Network object support for HTTP, ICMP, and SSH platform settings.	7.1.0	7.1.0	You can now use network object groups that contain network objects for hosts or networks when configuring the IP addresses in the Threat Defense Platform Settings policy. Supported platforms: Firewall Threat Defense
Support to specify trusted DNS servers.	7.1.0	7.1.0	The option to specify DNS servers that you can trust for address resolution while using direct internet access was introduced. We added a tab for configuring the trusted DNS servers when configuring direct internet access: **Devices** > **Platform Settings** > **DNS** > **Trusted DNS Servers**. Supported platforms: Firewall Threat Defense

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Platform settings using the MD5 authentication algorithm or DES encryption for SNMPv3 users can not be deployed to Firewall Threat Defense devices running Versions 7.0+.	7.0.0	7.0.0	The MD5 authentication algorithm and DES encryption for SNMPv3 users on Firewall Threat Defense were deprecated in Version 6.5. If your deployment includes SNMPv3 users using the MD5 authentication algorithm or DES encryption that were created using Version 6.4 or earlier, you can continue to use those users for Firewall Threat Defense devices running Version 6.7 or earlier. However, you cannot edit those users and retain the MD5 or DES settings, and you cannot create new users with the MD5 or DES settings. If your Firewall Management Center manages any Firewall Threat Defenses running Versions 7.0+, deploying a platform settings policy with SNMP v3 users using the MD5 authentication algorithm or DES encryption to those Firewall Threat Defenses will fail. New/modified screens: **Devices** > **Platform Settings** > **SNMP** > **Hosts** Supported platforms: Firewall Threat Defense
Specify SHA224 or SHA384 for SNMPv3 users' authorization algorithm.	7.0.0	7.0.0	You can now select SHA224 or SHA384 for SNMPv3 users' authorization algorithm. New/modified screens: **Devices** > **Platform Settings** > **SNMP** > **Users** Supported platforms: Firewall Threat Defense
Specify time zone for device.	6.6.0	6.6.0	Specify a local time zone for a managed device, for use in time-based policy application. New/modified screens: **Devices** > **Platform Settings** > **Time Zone** Supported platforms: Firewall Threat Defense
Specify the Management interface for SNMP communication.	6.6.0	6.6.0	You can now select the Management interface for communication between the device and the SNMP management station. New/modified screens: **Devices** > **Platform Settings** > **SNMP** > **Hosts** Supported platforms: Firewall Threat Defense
Specify SHA256 for SNMPv3 users' authorization algorithm.	6.6.0	6.6.0	You can now select SHA256 for SNMPv3 users' authorization algorithm. New/modified screens: **Devices** > **Platform Settings** > **SNMP** > **Users** Supported platforms: Firewall Threat Defense

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
DES encryption and the MD5 authentication algorithm for SNMPv3 users on Firewall Threat Defense have been deprecated.	6.5.0	Any	We recommend that you not use the MD5 authentication algorithm or DES encryption for SNMPv3 users on Firewall Threat Defense devices, as these options have been deprecated. If your deployment includes SNMPv3 users using the MD5 authentication algorithm or DES encryption that were created using a Version 6.4 or earlier, you can continue to use those users. However, you cannot edit those users and retain the MD5 or DES settings, and you cannot create new users with the MD5 or DES settings. New/modified screens: **Devices** > **Platform Settings** > **SNMP** > **Users** Supported platforms: Firewall Threat Defense
Allow user traffic to pass when TCP syslog server is down.	6.3.0	6.3.0	We recommend that you allow connections through the threat defense device when the external TCP syslog server is unreachable by the device. The **Allow user traffic to pass when TCP syslog server is down (Recommended to be enabled)** option in the Platform Settings is Enabled by default.
Limit number of SSH login failures.	6.3.0	6.3.0	When a user accesses any device via SSH and fails three successive login attempts, the device terminates the SSH session.
External authentication added for SSH.	6.2.3	6.2.3	You can now configure external authentication for SSH access to the Firewall Threat Defense using LDAP or RADIUS. New/modified screens: **Devices** > **Platform Settings** > **External Authentication** Supported platforms: Firewall Threat Defense
Support for UC/APPL compliance mode.	6.2.1	6.2.1	You can enable security certifications compliance in CC mode or UCAPL mode. Enabling security certifications compliance does not guarantee strict compliance with all requirements of the security mode selected. For more information on hardening procedures, refer to the guidelines for this product provided by the certifying entity. New/modified screens: **Devices** > **Platform Settings** > **UC/APPL Compliance** Supported platforms: Any device
SSL settings for remote access VPN.	6.2.1	6.2.1	The Firewall Threat Defense device uses the Secure Sockets Layer (SSL) protocol and Transport Layer Security (TLS) to support secure message transmission for Remote Access VPN connection from remote clients. You can configure SSL versions and encryption algorithms that will be negotiated and used for message transmission during remote VPN access over SSL. New/modified screens: **Devices** > **Platform Settings** > **SSL** Supported platforms: Firewall Threat Defense

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
External authentication for SSH and HTML removed.	6.1.0	6.1.0	Due to changes to support converged management access, only local users are supported for SSH and HTML to data interfaces. Also, you can no longer SSH to the logical Diagnostic interface; instead you can SSH to the logical Management interface (which shares the same physical port). Previously, only external authentication was supported for SSH and HTML access to Diagnostic and data interfaces, while only local users were supported to the Management interface. New/modified screens: **Devices** > **Platform Settings** > **External Authentication** Supported platforms: Firewall Threat Defense
Firepower Threat Defense support.	6.0.1	6.0.1	This feature was introduced. New/modified screens: **Devices** > **Platform Settings** Supported platforms: Firewall Threat Defense

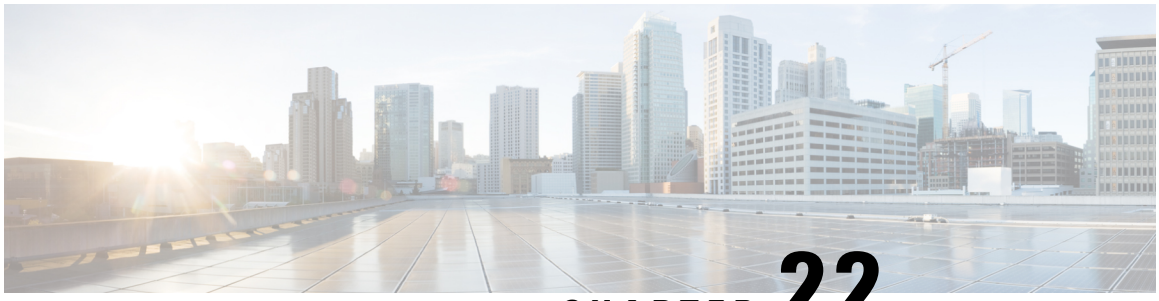

CHAPTER 22

Network Address Translation

The following topics explain Network Address Translation (NAT) and how to configure it on Firewall Threat Defense devices.

- Why Use NAT?, on page 1015
- NAT Basics, on page 1016
- Requirements and Prerequisites for NAT Policies, on page 1025
- Guidelines for NAT, on page 1025
- Manage NAT Policies, on page 1031
- Configure NAT for Threat Defense, on page 1033
- Translating IPv6 Networks, on page 1071
- Monitoring NAT, on page 1084
- Examples for NAT, on page 1085
- History for Firewall Threat Defense NAT, on page 1133

Why Use NAT?

Each computer and device within an IP network is assigned a unique IP address that identifies the host. Because of a shortage of public IPv4 addresses, most of these IP addresses are private, not routable anywhere outside of the private company network. RFC 1918 defines the private IP addresses you can use internally that should not be advertised:

- 10.0.0.0 through 10.255.255.255
- 172.16.0.0 through 172.31.255.255
- 192.168.0.0 through 192.168.255.255

One of the main functions of NAT is to enable private IP networks to connect to the Internet. NAT replaces a private IP address with a public IP address, translating the private addresses in the internal private network into legal, routable addresses that can be used on the public Internet. In this way, NAT conserves public addresses because it can be configured to advertise at a minimum only one public address for the entire network to the outside world.

Other functions of NAT include:

- Security—Keeping internal IP addresses hidden discourages direct attacks.
- IP routing solutions—Overlapping IP addresses are not a problem when you use NAT.

- Flexibility—You can change internal IP addressing schemes without affecting the public addresses available externally; for example, for a server accessible to the Internet, you can maintain a fixed IP address for Internet use, but internally, you can change the server address.
- Translating between IPv4 and IPv6 (Routed mode only) —If you want to connect an IPv6 network to an IPv4 network, NAT lets you translate between the two types of addresses.

Note NAT is not required. If you do not configure NAT for a given set of traffic, that traffic will not be translated, but will have all of the security policies applied as normal.

NAT Basics

The following topics explain some of the basics of NAT.

NAT Terminology

This document uses the following terminology:

- Real address/host/network/interface—The real address is the address that is defined on the host, before it is translated. In a typical NAT scenario where you want to translate the inside network when it accesses the outside, the inside network would be the "real" network. Note that you can translate any network connected to the device, not just an inside network. Therefore if you configure NAT to translate outside addresses, "real" can refer to the outside network when it accesses the inside network.
- Mapped address/host/network/interface—The mapped address is the address that the real address is translated to. In a typical NAT scenario where you want to translate the inside network when it accesses the outside, the outside network would be the "mapped" network.

Note During address translation, IP addresses configured for the device interfaces are not translated.

- Bidirectional initiation—Static NAT allows connections to be initiated *bidirectionally*, meaning both to the host and from the host.
- Source and destination NAT—For any given packet, both the source and destination IP addresses are compared to the NAT rules, and one or both can be translated/untranslated. For static NAT, the rule is bidirectional, so be aware that "source" and "destination" are used in commands and descriptions throughout this guide even though a given connection might originate at the "destination" address.

NAT Types

You can implement NAT using the following methods:

- Dynamic NAT—A group of real IP addresses are mapped to a (usually smaller) group of mapped IP addresses, on a first come, first served basis. Only the real host can initiate traffic. See Dynamic NAT, on page 1038.
- Dynamic Port Address Translation (PAT)—A group of real IP addresses are mapped to a single IP address using a unique source port of that IP address. See Dynamic PAT, on page 1043.
- Static NAT—A consistent mapping between a real and mapped IP address. Allows bidirectional traffic initiation. See Static NAT, on page 1053.
- Identity NAT—A real address is statically translated to itself, essentially bypassing NAT. You might want to configure NAT this way when you want to translate a large group of addresses, but then want to exempt a smaller subset of addresses. See Identity NAT, on page 1061.

NAT in Routed and Transparent Mode

You can configure NAT in both routed and transparent firewall mode. You cannot configure NAT for interfaces operating in inline, inline tap, or passive modes. The following sections describe typical usage for each firewall mode.

NAT in Routed Mode

The following figure shows a typical NAT example in routed mode, with a private network on the inside.

Figure 390: NAT Example: Routed Mode

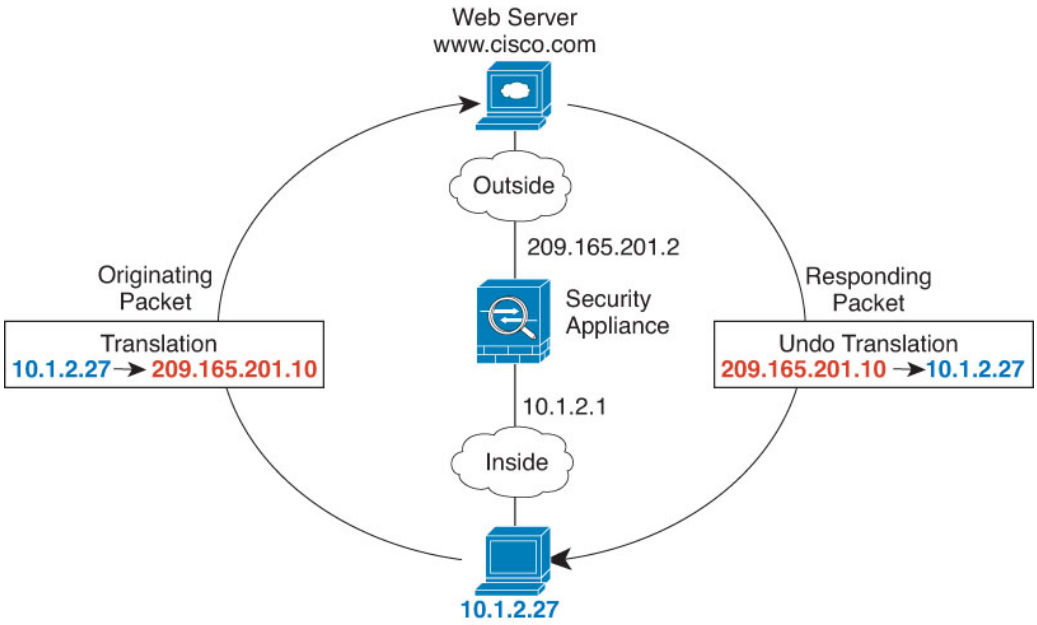

1. When the inside host at 10.1.2.27 sends a packet to a web server, the real source address of the packet, 10.1.2.27, is translated to a mapped address, 209.165.201.10.
2. When the server responds, it sends the response to the mapped address, 209.165.201.10, and the Firewall Threat Defense device receives the packet because the Firewall Threat Defense device performs proxy ARP to claim the packet.

3. The Firewall Threat Defense device then changes the translation of the mapped address, 209.165.201.10, back to the real address, 10.1.2.27, before sending it to the host.

NAT in Transparent Mode or Within a Bridge Group

Using NAT in transparent mode eliminates the need for the upstream or downstream routers to perform NAT for their networks. It can perform a similar function within a bridge group in routed mode.

NAT in transparent mode, or in routed mode between members of the same bridge group, has the following requirements and limitations:

- You cannot configure interface PAT when the mapped address is a bridge group member interface, because there is no IP address attached to the interface.
- ARP inspection is not supported. Moreover, if for some reason a host on one side of the Firewall Threat Defense sends an ARP request to a host on the other side of the Firewall Threat Defense, and the initiating host real address is mapped to a different address on the same subnet, then the real address remains visible in the ARP request.
- Translating between IPv4 and IPv6 networks is not supported. Translating between two IPv6 networks, or between two IPv4 networks is supported.

The following figure shows a typical NAT scenario in transparent mode, with the same network on the inside and outside interfaces. The transparent firewall in this scenario is performing the NAT service so that the upstream router does not have to perform NAT.

Figure 391: NAT Example: Transparent Mode

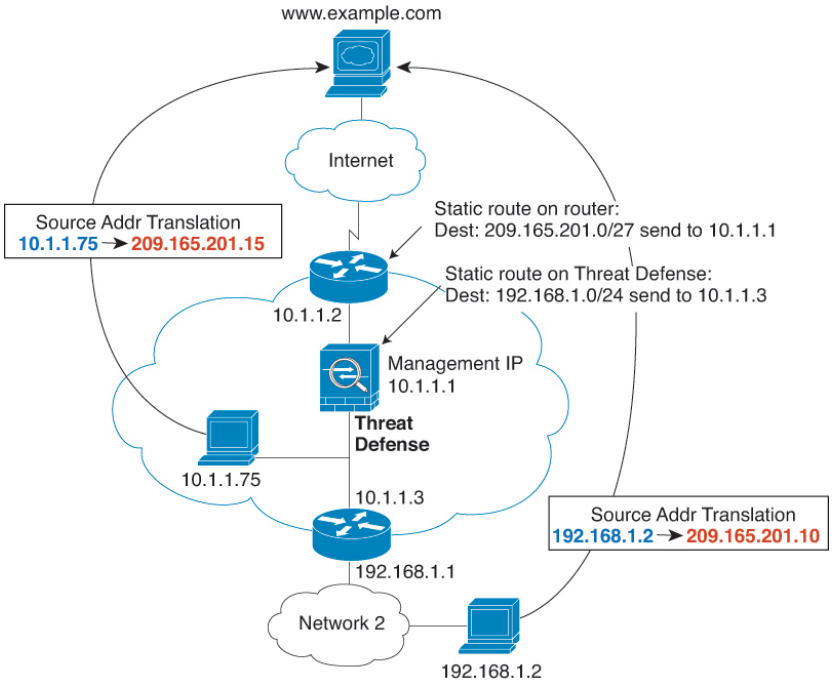

1. When the inside host at 10.1.1.75 sends a packet to a web server, the real source address of the packet, 10.1.1.75, is changed to a mapped address, 209.165.201.15.

2. When the server responds, it sends the response to the mapped address, 209.165.201.15, and the Firewall Threat Defense receives the packet because the upstream router includes this mapped network in a static route directed to the Firewall Threat Defense management IP address.
3. The Firewall Threat Defense then undoes the translation of the mapped address, 209.165.201.15, back to the real address, 10.1.1.1.75. Because the real address is directly-connected, the Firewall Threat Defense sends it directly to the host.
4. For host 192.168.1.2, the same process occurs, except for returning traffic, the Firewall Threat Defense looks up the route in its routing table and sends the packet to the downstream router at 10.1.1.3 based on the Firewall Threat Defense static route for 192.168.1.0/24.

Auto NAT and Manual NAT

You can implement address translation in two ways: *auto NAT* and *manual NAT*.

We recommend using auto NAT unless you need the extra features that manual NAT provides. It is easier to configure auto NAT, and it might be more reliable for applications such as Voice over IP (VoIP). (For VoIP, you might see a failure in the translation of indirect addresses that do not belong to either of the objects used in the rule.)

Auto NAT

All NAT rules that are configured as a parameter of a network object are considered to be *auto NAT* rules. This is a quick and easy way to configure NAT for a network object. You cannot create these rules for a group object, however.

Although these rules are configured as part of the object itself, you cannot see the NAT configuration in the object definition through the object manager.

When a packet enters an interface, both the source and destination IP addresses are checked against the auto NAT rules. The source and destination address in the packet can be translated by separate rules if separate matches are made. These rules are not tied to each other; different combinations of rules can be used depending on the traffic.

Because the rules are never paired, you cannot specify that sourceA/destinationA should have a different translation than sourceA/destinationB. Use manual NAT for that kind of functionality, where you can identify the source and destination address in a single rule.

Manual NAT

Manual NAT lets you identify both the source and destination address in a single rule. Specifying both the source and destination addresses lets you specify that sourceA/destinationA can have a different translation than sourceA/destinationB.

Note For static NAT, the rule is bidirectional, so be aware that “source” and “destination” are used in commands and descriptions throughout this guide even though a given connection might originate at the “destination” address. For example, if you configure static NAT with port address translation, and specify the source address as a Telnet server, and you want all traffic going to that Telnet server to have the port translated from 2323 to 23, then you must specify the *source* ports to be translated (real: 23, mapped: 2323). You specify the source ports because you specified the Telnet server address as the source address.

The destination address is optional. If you specify the destination address, you can either map it to itself (identity NAT), or you can map it to a different address. The destination mapping is always a static mapping.

Comparing Auto NAT and Manual NAT

The main differences between these two NAT types are:

- How you define the real address.
  - Auto NAT—The NAT rule becomes a parameter for a network object. The network object IP address serves as the original (real) address.
  - Manual NAT—You identify a network object or network object group for both the real and mapped addresses. In this case, NAT is not a parameter of the network object; the network object or group is a parameter of the NAT configuration. The ability to use a network object *group* for the real address means that manual NAT is more scalable.
- How source and destination NAT is implemented.
  - Auto NAT— Each rule can apply to either the source or destination of a packet. So two rules might be used, one for the source IP address, and one for the destination IP address. These two rules cannot be tied together to enforce a specific translation for a source/destination combination.
  - Manual NAT—A single rule translates both the source and destination. A packet matches one rule only, and further rules are not checked. Even if you do not configure the optional destination address, a matching packet still matches one manual NAT rule only. The source and destination are tied together, so you can enforce different translations depending on the source/destination combination. For example, sourceA/destinationA can have a different translation than sourceA/destinationB.
- Order of NAT Rules.
  - Auto NAT—Automatically ordered in the NAT table.
  - Manual NAT—Manually ordered in the NAT table (before or after auto NAT rules).

NAT Rule Order

Auto NAT and manual NAT rules are stored in a single table that is divided into three sections. Section 1 rules are applied first, then section 2, and finally section 3, until a match is found. For example, if a match is found in section 1, sections 2 and 3 are not evaluated. The following table shows the order of rules within each section.

Note There is also a Section 0, which contains any NAT rules that the system creates for its own use. These rules have priority over all others. The system automatically creates these rules and clears xlates as needed. You cannot add, edit, or modify rules in Section 0.

Table 68: NAT Rule Table

Table Section	Rule Type	Order of Rules within the Section
Section 1	Manual NAT	Applied on a first match basis, in the order they appear in the configuration. Because the first match is applied, you must ensure that specific rules come before more general rules, or the specific rules might not be applied as desired. By default, manual NAT rules are added to section 1. By "specific rules first," we mean: • Static rules should come before dynamic rules. • Rules that include destination translation should come before rules with source translation only. If you cannot eliminate overlapping rules, where more than one rule might apply based on the source or destination address, be especially careful to follow these recommendations.
Section 2	Auto NAT	If a match in section 1 is not found, section 2 rules are applied in the following order: **1.** Static rules. **2.** Dynamic rules. **Within each rule type, the following ordering guidelines are used:** **1.** Quantity of real IP addresses—From smallest to largest. For example, an object with one address will be assessed before an object with 10 addresses. **2.** For quantities that are the same, then the IP address number is used, from lowest to highest. For example, 10.1.1.0 is assessed before 11.1.1.0. **3.** If the same IP address is used, then the name of the network object is used, in alphabetical order. For example, abracadabra is assessed before catwoman.
Section 3	Manual NAT	If a match is still not found, section 3 rules are applied on a first match basis, in the order they appear in the configuration. This section should contain your most general rules. You must also ensure that any specific rules in this section come before general rules that would otherwise apply.

For section 2 rules, for example, you have the following IP addresses defined within network objects:

- 192.168.1.0/24 (static)
- 192.168.1.0/24 (dynamic)
- 10.1.1.0/24 (static)

- 192.168.1.1/32 (static)
- 172.16.1.0/24 (dynamic) (object def)
- 172.16.1.0/24 (dynamic) (object abc)

The resultant ordering would be:

- 192.168.1.1/32 (static)
- 10.1.1.0/24 (static)
- 192.168.1.0/24 (static)
- 172.16.1.0/24 (dynamic) (object abc)
- 172.16.1.0/24 (dynamic) (object def)
- 192.168.1.0/24 (dynamic)

NAT Interfaces

Except for bridge group member interfaces, you can configure a NAT rule to apply to any interface (in other words, all interfaces), or you can identify specific real and mapped interfaces. You can also specify any interface for the real address, and a specific interface for the mapped address, or vice versa.

For example, you might want to specify any interface for the real address and specify the outside interface for the mapped address if you use the same private addresses on multiple interfaces, and you want to translate them all to the same global pool when accessing the outside.

Figure 392: Specifying Any Interface

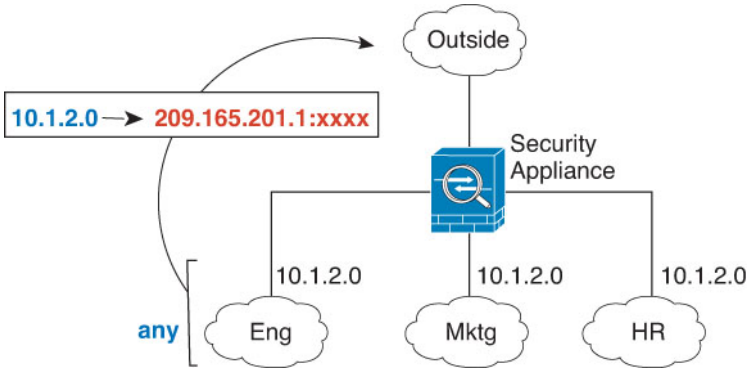

However, the concept of "any" interface does not apply to bridge group member interfaces. When you specify "any" interface, all bridge group member interfaces are excluded. Thus, to apply NAT to bridge group members, you must specify the member interface. This could result in many similar rules where only one interface is different. You cannot configure NAT for the Bridge Virtual Interface (BVI) itself, you can configure NAT for member interfaces only.

Note You cannot configure NAT for interfaces operating in inline, inline tap, or passive modes. When specifying interfaces, you do so indirectly by selecting the interface object that contains the interface.

NAT Exemption

When internet edge devices have a site-to-site VPN configured on an interface and also have NAT rules for that interface, you must exempt the VPN traffic from the NAT rules. If you do not exempt the VPN traffic from NAT translation, the traffic gets dropped or is not routed through the VPN tunnel to the remote peer.

NAT exemption allows you to exclude traffic from being translated by NAT rules. When you create a policy-based site-to-site VPN using the Firewall Management Center VPN wizard, you can select the **NAT Exempt** option to create the rules automatically (**Device > Site To Site**). By default, this option is enabled. You can view the NAT exemptions for a device in the NAT policy page (**Devices** > **NAT** , and then click **NAT Exemptions**).

The Firewall Management Center supports NAT exemption for all policy-based site-to-site VPN topology types. For more information, see Configure a Policy-based Site-to-Site VPN, on page 1516.

Consider the following example, which shows a site-to-site VPN tunnel connecting Site A and Site B. For traffic that must go to the Internet, NAT translates the private IPs to a public IP address to access the Internet. For traffic that must go over the VPN tunnel, you must configure NAT exemption for the device in the VPN wizard.

Figure 393: Site-to-site VPN Topology with NAT Exemption

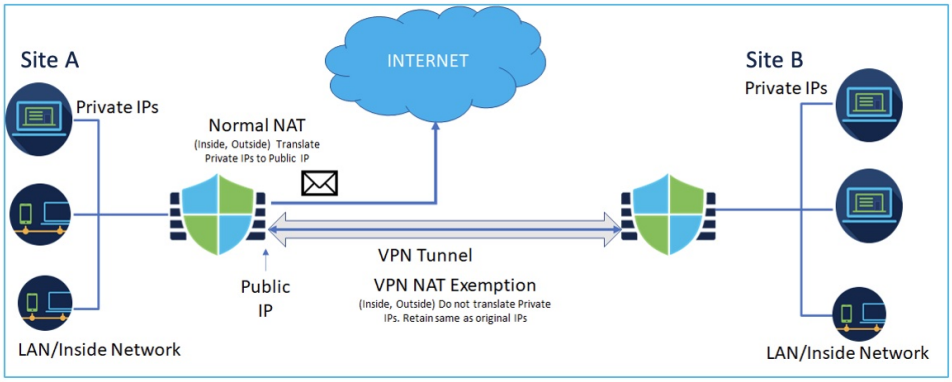

Note that when you upgrade from versions earlier than 7.4, if you have enabled NAT exemption, Firewall Management Center will disable this option. You must enable this option in the policy-based site-to-site VPN wizard.

Configuring Routing for NAT

The Firewall Threat Defense device needs to be the destination for any packets sent to the translated (mapped) address.

When sending packets, the device uses the destination interface if you specify one, or a routing table lookup if you do not, to determine the egress interface. For identity NAT, you have the option to use a route lookup even if you specify a destination interface.

The type of routing configuration needed depends on the type of mapped address, as explained in the following topics.

Addresses on the Same Network as the Mapped Interface

If you use addresses on the same network as the destination (mapped) interface, the Firewall Threat Defense device uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the Firewall Threat Defense device does not have to be the gateway for any additional networks. This solution is ideal if the outside network contains an adequate number of free addresses, a consideration if you are using a 1:1 translation like dynamic NAT or static NAT. Dynamic PAT greatly extends the number of translations you can use with a small number of addresses, so even if the available addresses on the outside network is small, this method can be used. For PAT, you can even use the IP address of the mapped interface.

Note If you configure the mapped interface to be any interface, and you specify a mapped address on the same network as one of the mapped interfaces, then if an ARP request for that mapped address comes in on a *different* interface, then you need to manually configure an ARP entry for that network on the ingress interface, specifying its MAC address. Typically, if you specify any interface for the mapped interface, then you use a unique network for the mapped addresses, so this situation would not occur. Configure the ARP table in the ingress interface's **Advanced** settings.

Addresses on a Unique Network

If you need more addresses than are available on the destination (mapped) interface network, you can identify addresses on a different subnet. The upstream router needs a static route for the mapped addresses that points to the Firewall Threat Defense device.

Alternatively for routed mode, you can configure a static route on the Firewall Threat Defense device for the mapped addresses using any IP address on the destination network as the gateway, and then redistribute the route using your routing protocol. For example, if you use NAT for the inside network (10.1.1.0/24) and use the mapped IP address 209.165.201.5, then you can configure a static route for 209.165.201.5 255.255.255.255 (host address) to the 10.1.1.99 gateway that can be redistributed.

For transparent mode, if the real host is directly-connected, configure the static route on the upstream router to point to the Firewall Threat Defense device: specify the bridge group IP address. For remote hosts in transparent mode, in the static route on the upstream router, you can alternatively specify the downstream router IP address.

The Same Address as the Real Address (Identity NAT)

The default behavior for identity NAT has proxy ARP enabled, matching other static NAT rules. You can disable proxy ARP if desired. You can also disable proxy ARP for regular static NAT if desired, in which case you need to be sure to have proper routes on the upstream router.

Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues. For example, if you configure a broad identity NAT rule for “any” IP address, then leaving proxy ARP enabled can cause problems for hosts on the network directly connected to the mapped interface. In this case, when a host on the mapped network wants to communicate with another host on the same network, then the address in the ARP request matches the NAT rule (which matches “any” address). The Firewall Threat Defense device will then proxy ARP for the address, even though the packet is not actually destined for the Firewall Threat Defense device. (Note that this problem occurs even if you have a manual NAT rule; although the NAT rule must match both the source and destination addresses, the proxy ARP decision is made only on the “source” address). If the Firewall Threat Defense device ARP response is received before the actual host ARP response, then traffic will be mistakenly sent to the Firewall Threat Defense device.

Requirements and Prerequisites for NAT Policies

Supported Domains

Any

User Roles

Admin

Access Admin

Network Admin

Guidelines for NAT

The following topics provide detailed guidelines for implementing NAT.

Firewall Mode Guidelines for NAT

NAT is supported in routed and transparent firewall mode.

However, configuring NAT on bridge group member interfaces (interfaces that are part of a Bridge Group Virtual Interface, or BVI) has the following restrictions:

- When configuring NAT for the members of a bridge group, you specify the member interface. You cannot configure NAT for the bridge group interface (BVI) itself.
- When doing NAT between bridge group member interfaces, you must specify the real and mapped addresses. You cannot specify “any” as the interface.
- You cannot configure interface PAT when the mapped address is a bridge group member interface, because there is no IP address attached to the interface.
- You cannot translate between IPv4 and IPv6 networks (NAT64/46) when the source and destination interfaces are members of the same bridge group. Static NAT/PAT 44/66, dynamic NAT44/66, and dynamic PAT44 are the only allowed methods; dynamic PAT66 is not supported. However, you can do NAT64/46 between members of different bridge groups, or between a bridge group member (source) and standard routed interface (destination).

Note You cannot configure NAT for interfaces operating in inline, inline tap, or passive modes.

IPv6 NAT Guidelines

NAT supports IPv6 with the following guidelines and restrictions.

- For standard routed mode interfaces, you can also translate between IPv4 and IPv6.

- You cannot translate between IPv4 and IPv6 for interfaces that are members of the same bridge group. You can translate between two IPv6 or two IPv4 networks only. This restriction does not apply when the interfaces are members of different bridge groups, or between a bridge group member and a standard routed interface.
- You cannot use dynamic PAT for IPv6 (NAT66) when translating between interfaces in the same bridge group. This restriction does not apply when the interfaces are members of different bridge groups, or between a bridge group member and a standard routed interface.
- For static NAT, you can specify an IPv6 subnet up to /64. Larger subnets are not supported.
- When using FTP with NAT46, when an IPv4 FTP client connects to an IPv6 FTP server, the client must use either the extended passive mode (EPSV) or extended port mode (EPRT); PASV and PORT commands are not supported with IPv6.

IPv6 NAT Best Practices

You can use NAT to translate between IPv6 networks, and also to translate between IPv4 and IPv6 networks (routed mode only). We recommend the following best practices:

- NAT66 (IPv6-to-IPv6)—We recommend using static NAT. Although you can use dynamic NAT or PAT, IPv6 addresses are in such large supply, you do not have to use dynamic NAT. If you do not want to allow returning traffic, you can make the static NAT rule unidirectional (manual NAT only).
- NAT46 (IPv4-to-IPv6)—We recommend using static NAT. Because the IPv6 address space is so much larger than the IPv4 address space, you can easily accommodate a static translation. If you do not want to allow returning traffic, you can make the static NAT rule unidirectional (manual NAT only). When translating to an IPv6 subnet (/96 or lower), the resulting mapped address is by default an IPv4-embedded IPv6 address, where the 32-bits of the IPv4 address is embedded after the IPv6 prefix. For example, if the IPv6 prefix is a /96 prefix, then the IPv4 address is appended in the last 32-bits of the address. For example, if you map 192.168.1.0/24 to 201b::0/96, then 192.168.1.4 will be mapped to 201b::0.192.168.1.4 (shown with mixed notation). If the prefix is smaller, such as /64, then the IPv4 address is appended after the prefix, and a suffix of 0s is appended after the IPv4 address. You can also optionally translate the addresses net-to-net, where the first IPv4 address maps to the first IPv6 address, the second to the second, and so on.
- NAT64 (IPv6-to-IPv4)—You may not have enough IPv4 addresses to accommodate the number of IPv6 addresses. We recommend using a dynamic PAT pool to provide a large number of IPv4 translations.

NAT Support for Inspected Protocols

Some application layer protocols that open secondary connections, or that embedded IP addresses in packets, are inspected to provide the following services:

- Pinhole creation—Some application protocols open secondary TCP or UDP connections either on standard or negotiated ports. Inspection opens pinholes for these secondary ports so that you do not need to create access control rules to allow them.
- NAT rewrite— Protocols such as FTP embed IP addresses and ports for the secondary connections in packet data as part of the protocol. If there is NAT translation involved for either of the endpoints, the inspection engines rewrite the packet data to reflect the NAT translation of the embedded addresses and ports. The secondary connections would not work without NAT rewrite.

- Protocol enforcement—Some inspections enforce some degree of conformance to the RFCs for the inspected protocol.

The following table lists the inspected protocols that apply NAT rewrite and their NAT limitations. Keep these limitations in mind when writing NAT rules that include these protocols. Inspected protocols not listed here do not apply NAT rewrite. These inspections include GTP, HTTP, IMAP, POP, SMTP, SSH, and SSL.

Note NAT rewrite is supported on the listed ports only. For some of these protocols, you can extend inspection to other ports using Network Analysis Policies, but NAT rewrite is not extended to those ports. This includes DCERPC, DNS, FTP, and Sun RPC inspection. If you use these protocols on non-standard ports, do not use NAT on the connections.

Table 69: NAT Supported Application Inspection

Application	Inspected Protocol, Port	NAT Limitations	Pinholes Created
DCERPC	TCP/135	No NAT64.	Yes
DNS over UDP	UDP/53	No NAT support is available for name resolution through WINS.	No
ESMTP	TCP/25	No NAT64.	No
FTP	TCP/21	(Clustering) No static PAT.	Yes
H.323 H.225 (Call signaling) H.323 RAS	TCP/1720 UDP/1718 For RAS, UDP/1718-1719	(Clustering) No static PAT. No extended PAT. No NAT64.	Yes
ICMP ICMP Error	ICMP (ICMP traffic directed to a device interface is never inspected.)	No limitations.	No
IP Options	RSVP	No NAT64.	No
NetBIOS Name Server over IP	UDP/137, 138 (Source ports)	No extended PAT. No NAT64.	No
RSH	TCP/514	No PAT. No NAT64. (Clustering) No static PAT.	Yes
RTSP	TCP/554 (No handling for HTTP cloaking.)	No extended PAT. No NAT64. (Clustering) No static PAT.	Yes

Application	Inspected Protocol, Port	NAT Limitations	Pinholes Created
SIP	TCP/5060 UDP/5060	No extended PAT. No NAT64 or NAT46. (Clustering) No static PAT.	Yes
Skinny (SCCP)	TCP/2000	No extended PAT. No NAT64, NAT46, or NAT66. (Clustering) No static PAT.	Yes
SQL*Net (versions 1, 2)	TCP/1521	No extended PAT. No NAT64. (Clustering) No static PAT.	Yes
Sun RPC	TCP/111 UDP/111	No extended PAT. No NAT64.	Yes
TFTP	UDP/69	No NAT64. (Clustering) No static PAT. Payload IP addresses are not translated.	Yes
XDMCP	UDP/177	No extended PAT. No NAT64. (Clustering) No static PAT.	Yes

FQDN Destination Guidelines

You can specify the translated (mapped) destination in a manual NAT rule using a fully-qualified domain name (FQDN) network object instead of an IP address. For example, you can create a rule based on traffic that is destined for the www.example.com web server.

When using an FQDN, the system obtains the DNS resolution and writes the NAT rule based on the returned address. If you are using multiple DNS server groups, the filter domains are honored and the address is requested from the appropriate group based on the filters. If more than one address is obtained from the DNS server, the address used is based on the following:

- If there is an address on the same subnet as the specified interface, that address is used. If there isn't one on the same subnet, the first address returned is used.
- The IP type for the translated source and translated destination must match. For example, if the translated source address is IPv6, the FQDN object must specify IPv6 as the address type. If the translated source is IPv4, the FQDN object can specify IPv4 or both IPv4 and IPv6. In this case, an IPv4 address is selected.

You cannot include an FQDN object in a network group that is used for manual NAT destination. In NAT, an FQDN object must be used alone, as only a single destination host makes sense for this type of NAT rule.

If the FQDN cannot be resolved to an IP address, the rule is not functional until a DNS resolution is obtained.

Additional Guidelines for NAT

- NAT rules apply to through the device traffic only, they do not apply to traffic initiated by the device, such as a RADIUS authentication.
- For interfaces that are members of a bridge group, you write NAT rules for the member interfaces. You cannot write NAT rules for the Bridge Virtual Interface (BVI) itself.
- You cannot write NAT rules for a Virtual Tunnel Interface (VTI), which are used in site-to-site VPN. Writing rules for the VTI's source interface will not apply NAT to the VPN tunnel. To write NAT rules that will apply to VPN traffic tunneled on a VTI, you must use "any" as the interface; you cannot explicitly specify interface names.
- (Auto NAT only.) You can only define a single NAT rule for a given object; if you want to configure multiple NAT rules for an object, you need to create multiple objects with different names that specify the same IP address.
- If a VPN is defined on an interface, inbound ESP traffic on the interface is not subject to the NAT rules. The system allows the ESP traffic for established VPN tunnels only, dropping traffic not associated with an existing tunnel. This restriction applies to ESP and UDP ports 500 and 4500.
- If you define a site-to-site VPN on a device that is behind a device that is applying dynamic PAT, so that UDP ports 500 and 4500 are not the ones actually used, you must initiate the connection from the device that is behind the PAT device. The responder cannot initiate the security association (SA) because it does not know the correct port numbers.
- If you change the NAT configuration, and you do not want to wait for existing translations to time out before the new NAT configuration is used, you can clear the translation table using the **clear xlate** command in the device CLI. However, clearing the translation table disconnects all current connections that use translations.

  If you create a new NAT rule that should apply to an existing connection (such as a VPN tunnel), you need to use **clear conn** to end the connection. Then, the attempt to re-establish the connection should hit the NAT rule and the connection should be NAT'ed correctly.

Note If you remove a dynamic NAT or PAT rule, and then add a new rule with mapped addresses that overlap the addresses in the removed rule, then the new rule will not be used until all connections associated with the removed rule time out or are cleared using the **clear xlate** or **clear conn** commands. This safeguard ensures that the same address is not assigned to multiple hosts.

- You cannot use an object group with both IPv4 and IPv6 addresses; the object group must include only one type of address.
- A network object used in NAT cannot include more than 131,838 IP addresses, either explicitly or implied in a range of addresses or a subnet. Break up the address space into smaller ranges and write separate rules for the smaller objects.
- (Manual NAT only.) When using **any** as the source address in a NAT rule, the definition of “any” traffic (IPv4 vs. IPv6) depends on the rule. Before the Firewall Threat Defense device performs NAT on a packet, the packet must be IPv6-to-IPv6 or IPv4-to-IPv4; with this prerequisite, the Firewall Threat Defense device can determine the value of **any** in a NAT rule. For example, if you configure a rule from “any” to an IPv6 server, and that server was mapped from an IPv4 address, then **any** means “any IPv6

traffic." If you configure a rule from "any" to "any," and you map the source to the interface IPv4 address, then **any** means "any IPv4 traffic" because the mapped interface address implies that the destination is also IPv4.

- You can use the same mapped object or group in multiple NAT rules.
- The mapped IP address pool cannot include:
  - The mapped interface IP address. If you specify "any" interface for the rule, then all interface IP addresses are disallowed. For interface PAT (routed mode only), specify the interface name instead of the interface address.
  - The failover interface IP address.
  - (Transparent mode.) The management IP address.
  - (Dynamic NAT.) The standby interface IP address when VPN is enabled.
- Avoid using overlapping addresses in static and dynamic NAT policies. For example, with overlapping addresses, a PPTP connection can fail to get established if the secondary connection for PPTP hits the static instead of dynamic xlate.
- You cannot use overlapping addresses in the source address of a NAT rule and a remote access VPN address pool.
- If you specify a destination interface in a rule, then that interface is used as the egress interface rather than looking up the route in the routing table. However, for identity NAT, you have the option to use a route lookup instead.
- If you use PAT on Sun RPC traffic, which is used to connect to NFS servers, be aware that the NFS server might reject connections if the PAT'ed port is above 1024. The default configuration of NFS servers is to reject connections from ports higher than 1024. The error is typically "Permission Denied." Mapping ports above 1024 happens if you do not select the option to include the reserved ports (1-1023) in the port range of a PAT pool. You can avoid this problem by changing the NFS server configuration to allow all port numbers.
- NAT applies to through traffic only. Traffic generated by the system is not subject to NAT.
- Do not name a network object or group pat-pool, using any combination of upper- or lower-case letters.
- The unidirectional option is mainly useful for testing purposes and might not work with all protocols. For example, SIP requires protocol inspection to translate SIP headers using NAT, but this will not occur if you make the translation unidirectional.
- You cannot use NAT on the internal payload of Protocol Independent Multicast (PIM) registers.
- (Manual NAT) When writing NAT rules for a dual ISP interface setup (primary and backup interfaces using service level agreements in the routing configuration), do not specify destination criteria in the rule. Ensure the rule for the primary interface comes before the rule for the backup interface. This allows the device to choose the correct NAT destination interface based on the current routing state when the primary ISP is unavailable. If you specify destination objects, the NAT rule will always select the primary interface for the otherwise duplicate rules.
- If you get the ASP drop reason nat-no-xlate-to-pat-pool for traffic that should not match the NAT rules defined for the interface, configure identity NAT rules for the affected traffic so the traffic can pass untranslated.

- If you configure NAT for GRE tunnel endpoints, you must disable keepalives on the endpoints or the tunnel cannot be established. The endpoints send keepalives to the original addresses.
- DHCP and BOOTP share ports UDP/67-68. Because BOOTP is obsolete, writing NAT rules for the bootps port can cause port allocation problems when also running DHCP. Consider using DHCP relay instead for transmitting DHCP requests between network segments.

Manage NAT Policies

Network Address Translation (NAT) converts the IP address of an incoming packet to a different address in the outgoing packet. One of the main functions of NAT is to enable private IP networks to connect to the Internet. NAT replaces a private IP address with a public IP address, translating the private addresses in the internal private network into routable addresses that can be used on the public Internet. NAT keeps track of the translations, also known as xlates, to ensure that return traffic is directed to the correct untranslated host address.

Procedure

Step 1 Choose **Devices** > **NAT** .

Step 2 Manage your NAT policies:

- Create—Click **New Policy**. See Creating NAT Policies, on page 1032.
- Copy—Click **Copy** ( ) next to the policy you want to copy. You are prompted to give the copy a new, unique name. The copy includes all policy rules and configurations, but does not include device assignments.
- Report—Click **Report** ( ) for the policy. You are prompted to save the PDF report, which includes policy attributes, device assignments, rules, and object usage information.
- Edit—Click **Edit** ( ) next to the policy you want to edit. See Configure NAT for Threat Defense, on page 1033.
- Delete—Click **Delete** ( ) next to the policy you want to delete, then click **OK**. When prompted whether to continue, you are also informed if another user has unsaved changes in the policy.

  Caution
  After you have deployed a NAT policy to a managed device, you cannot delete the policy from the device. Instead, you must deploy a NAT policy with no rules to remove the NAT rules already present on the managed device. You also cannot delete a policy that is the last deployed policy on any of its target devices, even if it is out of date. Before you can delete the policy completely, you must deploy a different policy to those targets.

Creating NAT Policies

When you create a new NAT policy you must, at minimum, give it a unique name. Although you are not required to identify policy targets at policy creation time, you must perform this step before you can deploy the policy. If you apply a NAT policy with no rules to a device, the system removes all NAT rules from that device.

Procedure

Step 1 Choose **Devices** > **NAT** .

Step 2 Click **New Policy**.

Step 3 Enter a unique **Name**.

Step 4 Optionally, enter a **Description**.

Step 5 Choose the devices where you want to deploy the policy:

- Choose a device in the **Available Devices** list, and click **Add to Policy**.
- Click and drag a device from the **Available Devices** list to the **Selected Devices** list.
- Remove a device from the **Selected Devices** list by clicking **Delete** ( ) next to the device.

Step 6 Click **Save**.

Configuring NAT Policy Targets

You can identify the managed devices you want to target with your policy while creating or editing a policy. You can search a list of available devices and high-availability pairs, and add them to a list of selected devices.

Procedure

Step 1 Choose **Devices** > **NAT** .

Step 2 Click **Edit** ( ) next to the NAT policy you want to modify.

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **Policy Assignments**.

Step 4 Do any of the following:

- To assign a device, high-availability pair, or device group to the policy, select it in the **Available Devices** list and click **Add to Policy**. You can also drag and drop.
- To remove a device assignment, click **Delete** ( ) next to a device, high-availability pair, or device group in the **Selected Devices** list.

Step 5 Click **OK**.

Configure NAT for Threat Defense

Network address translation can be very complex. We recommend that you keep your rules as simple as possible to avoid translation problems and difficult troubleshooting situations. Careful planning before you implement NAT is critical. The following procedure provides the basic approach.

The NAT policy is a shared policy. You assign the policy to devices that should have similar NAT rules.

Whether a given rule in the policy applies to an assigned device is determined by the interface objects (security zones or interface groups) used in the rule. If the interface objects include one or more interface for the device, the rule is deployed to the device. Thus, you can configure rules that apply to subsets of devices within a single shared policy by carefully designing your interface objects. Rules that apply to "any" interface object are deployed to all devices.

If you change the type of an interface to a type that is not valid for use with a NAT policy that targets a device with that interface, the policy labels the interface as deleted. Click **Save** in the NAT policy to automatically remove the interface from the policy.

You can configure multiple NAT policies if groups of your devices require significantly different rules.

Procedure

Step 1 Select **Devices** > **NAT** .

- Click **New Policy** to create a new policy. Give the policy a name, optionally assign devices to it, and click **Save**.

  You can change device assignments later by editing the policy and clicking **Policy Assignments**.

- Click **Edit** ( ) to edit an existing threat defense NAT policy.

  If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 2 Decide what kinds of rules you need.

You can create dynamic NAT, dynamic PAT, static NAT, and identity NAT rules. For an overview, see NAT Types, on page 1016.

Step 3 Decide which rules should be implemented as manual or auto NAT.

For a comparison of these two implementation options, see Auto NAT and Manual NAT, on page 1019.

Step 4 Decide which rules should be custom per device.

Because you can assign a NAT policy to multiple devices, you can configure a single rule on many devices. However, you might have rules that should be interpreted differently by each device, or some rules that should apply to a subset of devices only.

Use interface objects to control on which devices a rule is configured. Then, use object overrides on network objects to customize the addresses used per device.

For detailed information, see Customizing NAT Rules for Multiple Devices, on page 1034.

Step 5 Create the rules as explained in the following sections.

- Dynamic NAT, on page 1038
- Dynamic PAT, on page 1043
- Static NAT, on page 1053
- Identity NAT, on page 1061

Step 6 Manage the NAT policy and rules.

You can do the following to manage the policy and its rules.

- To edit the policy name or description, click in those fields, type in your changes, and click outside the fields.
- To view only those rules that apply to a specific device, click **Filter by Device** and select the desired device. A rule applies to a device if it uses an interface object that includes an interface on the device.
- To view any warnings or errors in the policy, click **Show Warnings**, then choose a **Device**. Warnings and errors mark configurations that could adversely affect traffic flow or prevent the policy from deploying.
- To change the devices to which the policy is assigned, click the **Policy Assignments** link and modify the selected devices list as desired.
- To change whether a rule is enabled or disabled, right click the rule and select the desired option from the **State** command. You can temporarily disable a rule without deleting it using these controls.
- To add a rule, click the **Add Rule** button.
- To edit a rule, click **Edit** (✎) for the rule.
- To delete a rule, click **Delete** (🗑) for the rule.
- To change the number of rules displayed on the page, use the **Rows Per Page** drop-down list.
- To select more than one rule to enable, disable, or delete, click the checkbox for the rules, or the checkbox in the header, then perform the action.

Step 7 Click **Save**.

You can now go to **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Customizing NAT Rules for Multiple Devices

Because the NAT policy is shared, you can assign a given policy to more than one device. However, you can configure at most one auto NAT rule for a given object. Thus, if you want to configure different translations for an object based on the specific device doing the translation, you need to carefully configure the interface objects (security zones or interface groups) and define network object overrides for the translated address.

The interface objects determine on which devices a rule gets configured. The network object overrides determine what IP addresses are used by a given device for that object.

Consider the following scenario:

- FTD-A and FTD-B have inside networks 192.168.1.0/24 attached to the interface named “inside.”
- On FTD-A, you want to translate all 192.168.1.0/24 addresses to a NAT pool in the 10.100.10.10 - 10.100.10.200 range when going to the “outside” interface.
- On FTD-B, you want to translate all 192.168.1.0/24 addresses to a NAT pool in the 10.200.10.10 - 10.200.10.200 range when going to the “outside” interface.

To accomplish the above, you would do the following. Although this example rule is for dynamic auto NAT, you can generalize the technique for any type of NAT rule.

Procedure

Step 1 Create the security zones for the inside and outside interfaces.

a) Choose **Objects** > **Object Management**.
b) Select **Interface Objects** from the table of contents and click **Add** > **Security Zone**. (You can use interface groups instead of zones.)
c) Configure the inside zone properties.
   - **Name**—Enter a name, for example, **inside-zone**.
   - **Type**—Select **Routed** for routed-mode devices, **Switched** for transparent mode.
   - **Selected Interfaces**—Add the FTD-A/inside and FTD-B/inside interfaces to the selected list.
d) Click **Save**.
e) Click **Add** > **Security Zone** and define the outside zone properties.
   - **Name**—Enter a name, for example, **outside-zone**.
   - **Interface Type**—Select **Routed** for routed-mode devices, **Switched** for transparent mode.
   - **Selected Interfaces**—Add the FTD-A/outside and FTD-B/outside interfaces to the selected list.
f) Click **Save**.

Step 2 Create the network object for the original inside network on the Object Management page.

a) Select **Network** from the table of contents and click **Add Network** > **Add Object**.
b) Configure the inside network properties.
   - **Name**—Enter a name, for example, **inside-network**.
   - **Network**—Enter the network address, for example, **192.168.1.0/24**.
c) Click **Save**.

Step 3 Create the network object for the translated NAT pool and define overrides.

a) Click **Add Network** > **Add Object**.
b) Configure the NAT pool properties for FTD-A.
   - **Name**—Enter a name, for example, **NAT-pool**.

- **Network**—Enter the range of addresses to include in the pool for FTD-A, for example, **10.100.10.10-10.100.10.200**.

c) Select **Allow Overrides**.
d) Click the **Overrides** heading to open the list of object overrides.
e) Click **Add** to open the Add Object Override dialog box.
f) Select FTD-B and **Add** it to the Selected Devices list.
g) Click **Override** and change **Network** to **10.200.10.10-10.200.10.200**
h) Click **Add** to add the override to the device.

By defining an override for FTD-B, whenever the system configures this object on FTD-B, it will use the override value instead of the value defined in the original object.

i) Click **Save**.

Step 4 Configure the NAT rule.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.
b) Click **Add Rule**.
c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Dynamic.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside-zone.
- **Destination Interface Objects** = outside-zone.

Note
The interface objects control on which devices the rule is configured. Because in this example the zones contain interfaces for FTD-A and FTD-B only, even if the NAT policy were assigned to additional devices, the rule would be deployed to those 2 devices only.

e) On **Translation**, configure the following:

- **Original Source** = inside-network object.
- **Translated Source** > **Address**= NAT-pool object.

f) Click **Save**.

You now have a single rule that will be interpreted differently for FTD-A and FTD-B, providing unique translations for the inside networks protected by each firewall.

Searching and Filtering the NAT Rule Table

You can search and filter the NAT rule table to help you find rules that you need to modify or view. When you filter the table, only matching rules are shown. Note that although the rule numbers change to be sequentially 1, 2, and so forth, filtering does not change the actual rule number or the rule's location in the

table relative to hidden rules. Filtering simply changes what you can see to help you locate rules that interest you.

When editing the NAT policy, you can use the fields above the table to do the following types of search/filter:

- **Filter by Device**—Click **Filter by Device**, then select the devices whose rules you want to see and click **OK**. Whether a rule applies to a device is determined by the rule's interface constraints. If you specify a security zone or interface group for either the source or destination interface, the rule applies to a device if at least one interface for the device is in the zone or group. If a NAT rule applies to any source and any destination interface, then it applies to all devices.

  If you also do a text or multiple-attribute search, the results are constrained to the selected devices.

  To remove this filter, click **Filter by Device** and deselect the devices, or select **All**, and click **OK**.

- **Simple Text Search**—In the **Filter** box, type a string and press Enter. The string is compared to all values in the rules. For example, if you enter "network-object-1," which is the name of a network object, you would get rules that use the object in source, destination, and PAT pool attributes.

  For network and port objects, the string is also compared to the contents of the objects used in the rule. For example, if a PAT pool object includes the range 10.100.10.3-10.100.10.100, searching on either 10.100.10.3 or 10.100.10.100 (or a partial 10.100.10) will include rules that use that PAT pool object. However, the match must be exact: searching on 10.100.10.5 will not match this PAT pool object, even though the IP address is within the object's IP address range.

  To remove the filter, click the **x** on the right side of the Filter box.

- **Multiple-Attribute Search**—If a simple text search gives you too many hits, you can configure multiple values for the search. Click in the **Filter** box to open the list of attributes, then select or enter strings for the attributes you intend to search and click the **Filter** button. These attributes are the same as the ones you would configure within a NAT rule. The attributes are AND'ed, so filtered results include only those rules that match all attributes you configured.

  - For binary attributes, such as the rule state (enabled/disabled), whether a PAT pool is configured (enabled/disabled), the direction of the rule (uni/bi), or rule type (static/dynamic), simply check or uncheck the boxes as appropriate. Select both boxes if you do not care about the attribute value. If you deselect both boxes, no rules will match the filter.
  - For string attributes, type a full or partial string relevant to that attribute. These will be object names, either for security zones/interface groups, network objects, or port objects. It can also be the network or port object contents, which are matched the same way they are for simple text searches.

  To remove the filter, click the **x** on the right side of the Filter box, or click in the Filter box to open the drop-down list, and click the Clear button.

Enabling, Disabling, or Deleting Multiple Rules

You can enable or disable manual NAT rules, or delete any NAT rule, one by one. You can also select multiple rules and apply changes to all of them at once. Because enable/disable applies to manual NAT only, if you select a mix of rule types, you can delete them only.

Note that when you enable or disable rules, it does not matter if you select some rules that were already enabled or disabled. For example, enabling an already enabled rule simply leaves the rule enabled.

Procedure

Step 1 Select **Devices** > **NAT** , and edit a threat defense NAT policy.

Step 2 (Optional.) Filter the NAT rules to locate the ones you want to change.

Filtering is especially useful if you have a large NAT policy. For example, you could search on disabled rules to find ones that need to be enabled.

Step 3 Select the rules you want to change.

- Click the checkbox in the left column of the rule to select (or deselect) individual rules.
- Click the checkbox in the table header to select all the rules on the currently displayed page.

Your selection is preserved as you go from page to page. However, in practice, it makes most sense to perform your actions on the rules selected on a page before going to the next page.

Step 4 Perform the desired action. When selecting multiple rules, you are asked to confirm the action.

Note that these actions are also available on the right-click menu.

- To enable all rules, click **Select Bulk Action** > **Enable**.
- To disable all rules, click **Select Bulk Action** > **Disable**.
- To delete all rules, click **Select Bulk Action** > **Delete**.

Dynamic NAT

The following topics explain dynamic NAT and how to configure it.

About Dynamic NAT

Dynamic NAT translates a group of real addresses to a pool of mapped addresses that are routable on the destination network. The mapped pool typically includes fewer addresses than the real group. When a host you want to translate accesses the destination network, NAT assigns the host an IP address from the mapped pool. The translation is created only when the real host initiates the connection. The translation is in place only for the duration of the connection, and a given user does not keep the same IP address after the translation times out. Users on the destination network, therefore, cannot initiate a reliable connection to a host that uses dynamic NAT, even if the connection is allowed by an access rule.

Note For the duration of the translation, a remote host can initiate a connection to the translated host if an access rule allows it. Because the address is unpredictable, a connection to the host is unlikely. Nevertheless, in this case you can rely on the security of the access rule. A successful connection from a remote host can reset the idle timer for the connection.

The following figure shows a typical dynamic NAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back.

Figure 394: Dynamic NAT

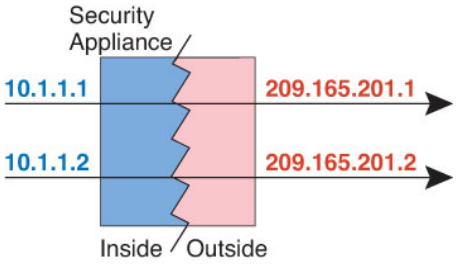

The following figure shows a remote host attempting to initiate a connection to a mapped address. This address is not currently in the translation table; therefore, the packet is dropped.

Figure 395: Remote Host Attempts to Initiate a Connection to a Mapped Address

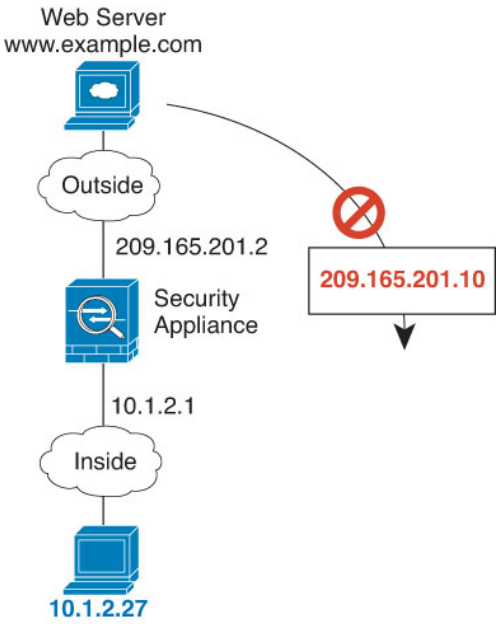

Dynamic NAT Disadvantages and Advantages

Dynamic NAT has these disadvantages:

- If the mapped pool has fewer addresses than the real group, you could run out of addresses if the amount of traffic is more than expected.

  Use PAT or a PAT fall-back method if this event occurs often because PAT provides over 64,000 translations using ports of a single address.

- You have to use a large number of routable addresses in the mapped pool, and routable addresses may not be available in large quantities.

The advantage of dynamic NAT is that some protocols cannot use PAT. PAT does not work with the following:

- IP protocols that do not have a port to overload, such as GRE version 0.
- Some multimedia applications that have a data stream on one port, the control path on another port, and are not open standard.

Configure Dynamic Auto NAT

Use dynamic auto NAT rules to translate addresses to different IP addresses that are routable on the destination network.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Source**—This must be a network object (not a group), and it can be a host, range, or subnet.
- **Translated Source**—This can be a network object or group, but it cannot include a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. If a group contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Auto NAT Rule**.
- **Type**—Select **Dynamic**.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 On **Translation**, configure the following options:

- **Original Source**—The network object that contains the addresses you are translating.
- **Translated Source**—The network object or group that contains the mapped addresses.

Step 6 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see Rewriting DNS Queries and Responses Using NAT, on page 1117.
- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.

Step 7 Click **Save** to add the rule.

Step 8 Click **Save** on the NAT page to save your changes.

Configure Dynamic Manual NAT

Use dynamic manual NAT rules when auto NAT does not meet your needs. For example, if you want to do different translations based on the destination. Dynamic NAT translates addresses to different IP addresses that are routable on the destination network.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source**—This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source**—This can be a network object or group, but it cannot include a subnet. If a group contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback.

You can also create network objects or groups for the **Original Destination** and **Translated Destination** if you are configuring a static translation for those addresses in the rule.

For dynamic NAT, you can also perform port translation on the destination. In the Object Manager, ensure that there are port objects you can use for the **Original Destination Port** and **Translated Destination Port**. If you specify the source port, it will be ignored.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Manual NAT Rule**.
- **Type**—Select **Dynamic**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- **Enable**—Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page.
- **Insert**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 (On the **Translation** page.) Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet.

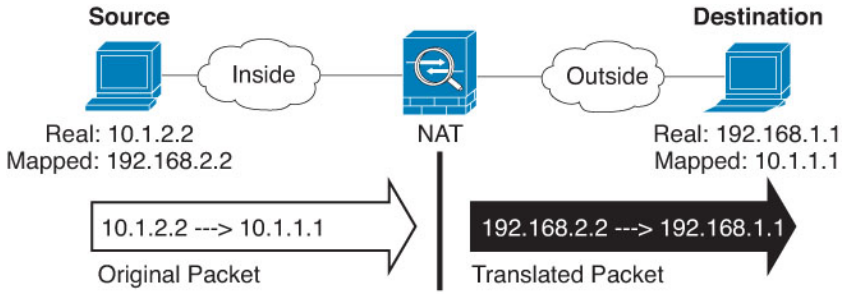

- **Original Source**—The network object or group that contains the addresses you are translating.
- **Original Destination**—(Optional.) The network object or group that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

  You can select **Source Interface IP** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source**—The network object or group that contains the mapped addresses.
- **Translated Destination**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the destination service ports for service translation: **Original Destination Port**, **Translated Destination Port**.

Dynamic NAT does not support port translation, so leave the **Original Source Port** and **Translated Source Port** fields empty. However, because the destination translation is always static, you can perform port translation for the destination port.

NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

Step 8 (Optional.) On **Advanced**, select the desired options:

- (For source translation only.) **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see Rewriting DNS Queries and Responses Using NAT, on page 1117.
- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.

Step 9 Click **Save** to add the rule.

Step 10 Click **Save** on the NAT page to save your changes.

Dynamic PAT

The following topics describe dynamic PAT.

About Dynamic PAT

Dynamic PAT translates multiple real addresses to a single mapped IP address by translating the real address and source port to the mapped address and a unique port.

Each connection requires a separate translation session because the source port differs for each connection. For example, 10.1.1.1:1025 requires a separate translation from 10.1.1.1:1026.

The following figure shows a typical dynamic PAT scenario. Only real hosts can create a NAT session, and responding traffic is allowed back. The mapped address is the same for each translation, but the port is dynamically assigned.

Figure 396: Dynamic PAT

For the duration of the translation, a remote host on the destination network can initiate a connection to the translated host if an access rule allows it. Because the port address (both real and mapped) is unpredictable, a connection to the host is unlikely. Nevertheless, in this case you can rely on the security of the access rule.

After the connection expires, the port translation also expires.

Note We recommend that you use different PAT pools for each interface. If you use the same pool for multiple interfaces, especially if you use it for "any" interface, the pool can be quickly exhausted, with no ports available for new translations.

Dynamic PAT Disadvantages and Advantages

Dynamic PAT lets you use a single mapped address, thus conserving routable addresses. You can even use the Firewall Threat Defense device interface IP address as the PAT address.

You cannot use dynamic PAT for IPv6 (NAT66) when translating between interfaces in the same bridge group. This restriction does not apply when the interfaces are members of different bridge groups, or between a bridge group member and a standard routed interface.

Dynamic PAT does not work with some multimedia applications that have a data stream that is different from the control path. For more information, see NAT Support for Inspected Protocols, on page 1026.

Dynamic PAT might also create a large number of connections appearing to come from a single IP address, and servers might interpret the traffic as a DoS attack. You can configure a PAT pool of addresses and use a round-robin assignment of PAT addresses to mitigate this situation.

PAT Pool Object Guidelines

When creating network objects for a PAT pool, follow these guidelines.

For a PAT pool

- Ports are mapped to an available port in the 1024 to 65535 range. You can optionally include the reserved ports, those below 1024, to make the entire port range available for translations.

  When operating in a cluster, blocks of 512 ports per address are allocated to the members of the cluster, and mappings are made within these port blocks. If you also enable block allocation, the ports are distributed according to the block allocation size, whose default is also 512. If you change the cluster unit limit (the size of the cluster), ensure that you clear xlates or reboot the devices so that PAT pools can be appropriately reallocated to the cluster units.

- If you enable block allocation for a PAT pool, port blocks are allocated in the 1024-65535 range only. Thus, if an application requires a low port number (1-1023), it might not work. For example, an application requesting port 22 (SSH) will get a mapped port within the range of 1024-65535 and within the block allocated to the host.
- If you use the same PAT pool object in two separate rules, then be sure to specify the same options for each rule. For example, if one rule specifies extended PAT, then the other rule must also specify extended PAT.
- If a host has an existing connection, then subsequent connections from that host use the same PAT IP address. If no ports are available, this can prevent the connection. Use the round robin option to avoid this problem.
- For best performance, limit the number of IP addresses within a PAT pool to 10,000.

For extended PAT for a PAT pool

- Many application inspections do not support extended PAT.
- If you enable extended PAT for a dynamic PAT rule, then you cannot also use an address in the PAT pool as the PAT address in a separate static NAT with port translation rule. For example, if the PAT pool includes 10.1.1.1, then you cannot create a static NAT-with-port-translation rule using 10.1.1.1 as the PAT address.
- If you use a PAT pool and specify an interface for fallback, you cannot specify extended PAT.
- For VoIP deployments that use ICE or TURN, do not use extended PAT. ICE and TURN rely on the PAT binding to be the same for all destinations.
- You cannot use extended PAT on units in a cluster.
- Extended PAT increases memory usage on the device.

For round robin for a PAT pool

- If a host has an existing connection, then subsequent connections from that host will use the same PAT IP address if ports are available. However, this “stickiness” does not survive a failover. If the device fails over, then subsequent connections from a host might not use the initial IP address.
- IP address “stickiness” is also impacted if you mix PAT pool/round robin rules with interface PAT rules on the same interface. For any given interface, choose either a PAT pool or interface PAT; do not create competing PAT rules.
- Round robin, especially when combined with extended PAT, can consume a large amount of memory. Because NAT pools are created for every mapped protocol/IP address/port range, round robin results in a large number of concurrent NAT pools, which use memory. Extended PAT results in an even larger number of concurrent NAT pools.

Configure Dynamic Auto PAT

Use dynamic auto PAT rules to translate addresses to unique IP address/port combinations, rather than to multiple IP addresses only. You can translate to a single address (either the destination interface's address or another address), or use a PAT pool of addresses to provide a larger number of possible translations.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Source**—This must be a network object (not a group), and it can be a host, range, or subnet.
- **Translated Source**—You have the following options to specify the PAT address:
  - **Destination Interface**—To use the destination interface address, you do not need a network object.
  - **Single PAT address**—Create a network object containing a single host.
  - **PAT pool**—Create a network object that includes a range, or create a network object group that contains hosts, ranges, or both. You cannot include subnets. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Auto NAT Rule**.
- **Type**—Select **Dynamic**.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 On **Translation**, configure the following options:

- **Original Source**—The network object that contains the addresses you are translating.
- **Translated Source**—One of the following:
  - (Interface PAT.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. Skip the step for configuring a PAT pool.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose. Skip the step for configuring a PAT pool.
  - To use a PAT pool, leave **Translated Source** empty.

Step 6 If you are using a PAT pool, select the **PAT Pool** page and do the following:

a) Select **Enable PAT pool**.

b) Select the network object group that contains the addresses for the pool in the **PAT** > **Address** field.

You can alternatively select **Destination Interface IP**, which is another way to implement interface PAT.

c) (Optional) Select the following options as needed:

- **Use Round Robin Allocation**—To assign addresses/ports in a round-robin fashion. By default without round robin, all ports for a PAT address will be allocated before the next PAT address is used. The round-robin method assigns one address/port from each PAT address in the pool before returning to use the first address again, and then the second address, and so on.
- **Extended PAT Table**—To use extended PAT. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80. You cannot use this option with interface PAT or interface PAT fallback.
- **Flat Port Range**, **Include Reserved Ports**—To use the 1024 to 65535 port range as a single flat range when allocating TCP/UDP ports. (Pre-6.7) When choosing the mapped port number for a translation, PAT uses the real source port number if it is available. However, without this option, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 1 to 511, 512 to 1023, and 1024 to 65535. To avoid running out of ports at the low ranges, configure this setting. To use the entire range of 1 to 65535, also check the **Include Reserved Ports** option. For the Firewall Threat Defense devices running version 6.7 or higher, the flat port range is always configured, whether you select the option or not. You can still select the **Include Reserved Ports** option for these systems, and that setting is honored.
- **Block Allocation**—To enable port block allocation. For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time. If you allocate a block of ports, subsequent connections from the host use new randomly selected ports within the block. If necessary, additional blocks are allocated if the host has active connections for all ports in the original block. Port blocks are allocated in the 1024-65535 range only. Port block allocation is compatible with round robin, but you cannot use it with the extended PAT or flat port range options. You also cannot use interface PAT fallback.

Step 7 (Optional.) On **Advanced**, select the desired options:

- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option. You cannot select this option if you already configured interface PAT as the translated address or PAT pool.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.

Step 8 Click **Save** to add the rule.

Step 9 Click **Save** on the NAT page to save your changes.

Configure Dynamic Manual PAT

Use dynamic manual PAT rules when auto PAT does not meet your needs. For example, if you want to do different translations based on the destination. Dynamic PAT translates addresses to unique IP address/port combinations, rather than to multiple IP addresses only. You can translate to a single address (either the destination interface's address or another address), or use a PAT pool of addresses to provide a larger number of possible translations.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source**—This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source**—You have the following options to specify the PAT address:
  - **Destination Interface**—To use the destination interface address, you do not need a network object.
  - **Single PAT address**—Create a network object containing a single host.
  - **PAT pool**—Create a network object that includes a range, or create a network object group that contains hosts, ranges, or both. You cannot include subnets.

You can also create network objects or groups for the **Original Destination** and **Translated Destination** if you are configuring a static translation for those addresses in the rule.

For dynamic NAT, you can also perform port translation on the destination. In the Object Manager, ensure that there are port objects you can use for the **Original Destination Port** and **Translated Destination Port**. If you specify the source port, it will be ignored.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Manual NAT Rule**.
- **Type**—Select **Dynamic**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- **Enable**—Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page.
- **Insert**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 (On the **Translation** page.) Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet.

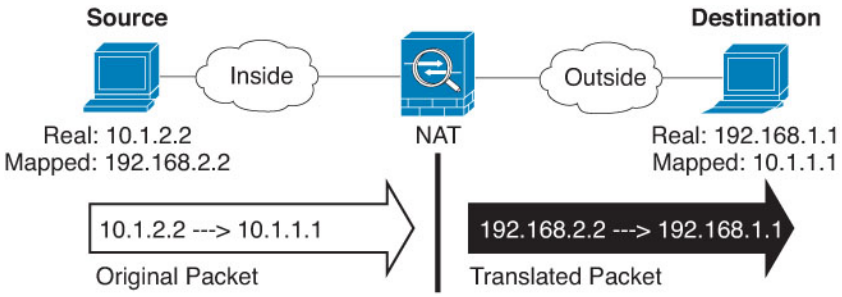

- **Original Source**—The network object or group that contains the addresses you are translating.
- **Original Destination**—(Optional.) The network object or group that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

  You can select **Source Interface IP** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source**—One of the following:
  - (Interface PAT.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. Skip the step for configuring a PAT pool.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose. Skip the step for configuring a PAT pool.
  - To use a PAT pool, leave **Translated Source** empty.
- **Translated Destination**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the destination service ports for service translation: **Original Destination Port**, **Translated Destination Port**.

Dynamic NAT does not support port translation, so leave the **Original Source Port** and **Translated Source Port** fields empty. However, because the destination translation is always static, you can perform port translation for the destination port.

NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

Step 8 If you are using a PAT pool, select the **PAT Pool** page and do the following:

a) Select **Enable PAT pool**.

b) Select the network object group that contains the addresses for the pool in the **PAT** > **Address** field.

You can alternatively select **Destination Interface IP**, which is another way to implement interface PAT.

c) (Optional) Select the following options as needed:

- **Use Round Robin Allocation**—To assign addresses/ports in a round-robin fashion. By default without round robin, all ports for a PAT address will be allocated before the next PAT address is used. The round-robin method assigns one address/port from each PAT address in the pool before returning to use the first address again, and then the second address, and so on.
- **Extended PAT Table**—To use extended PAT. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80. You cannot use this option with interface PAT or interface PAT fallback.
- **Flat Port Range**, **Include Reserved Ports**—To use the 1024 to 65535 port range as a single flat range when allocating TCP/UDP ports. (Pre-6.7) When choosing the mapped port number for a translation, PAT uses the real source port number if it is available. However, without this option, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 1 to 511, 512 to 1023, and 1024 to 65535. To avoid running out of ports at the low ranges, configure this setting. To use the entire range of 1 to 65535, also check the **Include Reserved Ports** option. For the Firewall Threat Defense devices running version 6.7 or higher, the flat port range is always configured, whether you select the option or not. You can still select the **Include Reserved Ports** option for these systems, and that setting is honored.
- **Block Allocation**—To enable port block allocation. For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time. If you allocate a block of ports, subsequent connections from the host use new randomly selected ports within the block. If necessary, additional blocks are allocated if the host has active connections for all ports in the original block. Port blocks are allocated in the 1024-65535 range only. Port block allocation is compatible with round robin, but you cannot use it with the extended PAT or flat port range options. You also cannot use interface PAT fallback.

Step 9 (Optional.) On **Advanced**, select the desired options:

- **Fallthrough to Interface PAT (Destination Interface)**—Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.

Step 10 Click **Save** to add the rule.

Step 11 Click **Save** on the NAT page to save your changes.

Configure PAT with Port Block Allocation

For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time (see RFC 6888). If you allocate a block of ports, subsequent connections from the host use new randomly-selected ports within the block. If necessary, additional blocks are allocated if the host has active connections for all ports in the original block. Blocks are freed when the last xlate that uses a port in the block is removed.

The main reason for allocating port blocks is reduced logging. The port block allocation is logged, connections are logged, but xlates created within the port block are not logged. On the other hand, this makes log analysis more difficult.

Port blocks are allocated in the 1024-65535 range only. There are separate blocks for TCP, UDP, and ICMP connections, and these blocks can overlap. Thus, if an application requires a low port number (1-1023), it might not work. For example, an application requesting port 22 (SSH) will get a mapped port within the range of 1024-65535 and within the block allocated to the host. You can create a separate NAT rule that does not use block allocation for applications that use low port numbers; for twice NAT, ensure the rule comes before the block allocation rule.

Before you begin

Usage notes for NAT rules:

- You can include the **Use Round Robin Allocation** option, but you cannot include the options for extending PAT uniqueness, using a flat range, including the reserved ports, or falling through to interface PAT. Other source/destination address and port information is also allowed.
- As with all NAT changes, if you replace an existing rule, you must clear xlates related to the replaced rule to have the new rule take effect. You can clear them explicitly or simply wait for them to time out. However, when operating in a cluster, you must reload the cluster so that the units can distribute blocks correctly.

Note If you are switching between a regular PAT and block allocation PAT rule, for object NAT, you must first delete the rule, then clear xlates. You can then create the new object NAT rule. Otherwise, you will see pat-port-block-state-mismatch drops in the **show asp drop** output.

- For a given PAT pool, you must specify (or not specify) block allocation for all rules that use the pool. You cannot allocate blocks in one rule and not in another. PAT pools that overlap also cannot mix block allocation settings. You also cannot overlap static NAT with port translation rules with the pool.

Procedure

Step 1 (Optional.) Configure global PAT port block allocation settings.

There are a few global settings that control port block allocation. If you want to change the defaults for these options, you must configure a FlexConfig object and add it to your FlexConfig policy.

a) Select **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object** and create a new object.
b) Configure the block allocation size, which is the number of ports in each block.

   xlate block-allocation size *value*

   The range is 32-4096. The default is 512. Use the "no" form to return to the default.

   If you do not use the default, ensure that the size you choose divides evenly into 64,512 (the number of ports in the 1024-65535 range). Otherwise, there will be ports that cannot be used. For example, if you specify 100, there will be 12 unused ports.

c) Configure the maximum blocks that can be allocated per host.

   xlate block-allocation maximum-per-host *number*

   The limit is per protocol, so a limit of 4 means at most 4 UDP blocks, 4 TCP blocks, and 4 ICMP blocks per host. The range is 1-8, the default is 4. Use the "no" form to return to the default.

d) (Optional.) Enable interim syslog generation.

   xlate block-allocation pba-interim-logging *seconds*

   By default, the system generates syslog messages during port block creation and deletion. If you enable interim logging, the system generates the following message at the interval you specify. The messages report all active port blocks allocated at that time, including the protocol (ICMP, TCP, UDP) and source and destination interface and IP address, and the port block. You can specify an interval from 21600-604800 seconds (6 hours to 7 days).

   %ASA-6-305017: Pba-interim-logging: Active *protocol* block of ports for translation from *real_interface*:*real_host_ip* to *mapped_interface*:*mapped_ip_address*/*start_port_num-end_port_num*

   Example:

   The following example sets the block allocation size to 64, the per-host maximum to 8, and enables interim logging every 6 hours.

   ```
   xlate block-allocation size 64
   xlate block-allocation maximum-per-host 8
   xlate block-allocation pba-interim-logging 21600
   ```

e) Select the following options in the FlexConfig object:

   - **Deployment** = **Everytime**
   - **Type** = **Append**

f) Click **Save** to create the FlexConfig object.
g) Select **Devices** > **FlexConfig**, and create or edit the FlexConfig policy that is assigned to the devices that need to have these settings adjusted.
h) Select your object in the available objects list and click > to move it to the selected objects list.
i) Click **Save**.

   You can click **Preview Config**, select one of the target devices, and verify that the xlate commands appear correctly.

Step 2 Add NAT rules that use PAT pool port block allocation.

a) Select **Devices** > **NAT** and add or edit the threat defense NAT policy.
b) Add or edit a NAT rule and configure at least the following options.
   - **Type** = **Dynamic**
   - In **Translation** > **Original Source**, select the object that defines the source address.
   - On **PAT Pool**, configure the following options:
     - Select **Enable PAT Pool**
     - In **PAT** > **Address**, select a network object or group that defines the pat pool.
     - Select the **Block Allocation** option.
c) Save your changes to the rule and to the NAT policy.

Static NAT

The following topics explain static NAT and how to implement it.

About Static NAT

Static NAT creates a fixed translation of a real address to a mapped address. Because the mapped address is the same for each consecutive connection, static NAT allows bidirectional connection initiation, both to and from the host (if an access rule exists that allows it). With dynamic NAT and PAT, on the other hand, each host uses a different address or port for each subsequent translation, so bidirectional initiation is not supported.

The following figure shows a typical static NAT scenario. The translation is always active so both real and remote hosts can initiate connections.

Figure 397: Static NAT

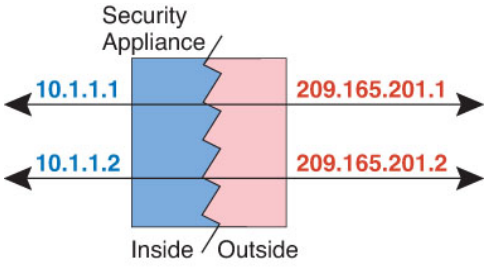

Note You can disable bidirectionality if desired.

Static NAT with Port Translation

Static NAT with port translation lets you specify a real and mapped protocol and port.

When you specify the port with static NAT, you can choose to map the port and/or the IP address to the same value or to a different value.

The following figure shows a typical static NAT with port translation scenario showing both a port that is mapped to itself and a port that is mapped to a different value; the IP address is mapped to a different value in both cases. The translation is always active so both translated and remote hosts can initiate connections.

Figure 398: Typical Static NAT with Port Translation Scenario

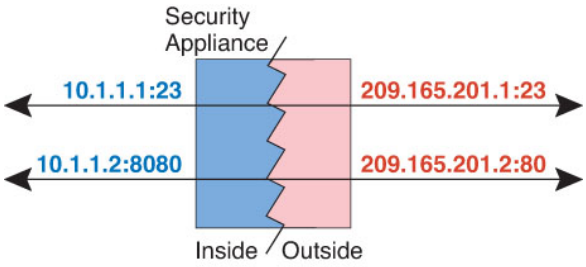

Static NAT-with-port-translation rules limit access to the destination IP address for the specified port only. If you try to access the destination IP address on a different port not covered by a NAT rule, then the connection is blocked. In addition, for manual NAT, traffic that does not match the source IP address of the NAT rule will be dropped if it matches the destination IP address, regardless of the destination port. Therefore, you must add additional rules for all other traffic allowed to the destination IP address. For example, you can configure a static NAT rule for the IP address, without port specification, and place it after the port translation rule.

Note For applications that require application inspection for secondary channels (for example, FTP and VoIP), NAT automatically translates the secondary ports.

Following are some other uses of static NAT with port translation.

Static NAT with Identity Port Translation

You can simplify external access to internal resources. For example, if you have three separate servers that provide services on different ports (such as FTP, HTTP, and SMTP), you can give external users a single IP address to access those services. You can then configure static NAT with identity port translation to map the single external IP address to the correct IP addresses of the real servers based on the port they are trying to access. You do not need to change the port, because the servers are using the standard ones (21, 80, and 25 respectively).

Static NAT with Port Translation for Non-Standard Ports

You can also use static NAT with port translation to translate a well-known port to a non-standard port or vice versa. For example, if inside web servers use port 8080, you can allow outside users to connect to port 80, and then undo translation to the original port 8080. Similarly, to provide extra security, you can tell web users to connect to non-standard port 6785, and then undo translation to port 80.

Static Interface NAT with Port Translation

You can configure static NAT to map a real address to an interface address/port combination. For example, if you want to redirect Telnet access for the device's outside interface to an inside host, then you can map the inside host IP address/port 23 to the outside interface address/port 23.

One-to-Many Static NAT

Typically, you configure static NAT with a one-to-one mapping. However, in some cases, you might want to configure a single real address to several mapped addresses (one-to-many). When you configure one-to-many

static NAT, when the real host initiates traffic, it always uses the first mapped address. However, for traffic initiated to the host, you can initiate traffic to any of the mapped addresses, and they will be untranslated to the single real address.

The following figure shows a typical one-to-many static NAT scenario. Because initiation by the real host always uses the first mapped address, the translation of real host IP/first mapped IP is technically the only bidirectional translation.

Figure 399: One-to-Many Static NAT

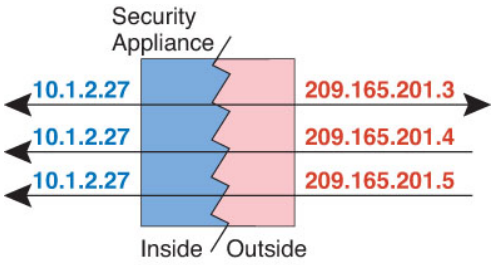

For example, you have a load balancer at 10.1.2.27. Depending on the URL requested, it redirects traffic to the correct web server.

Figure 400: One-to-Many Static NAT Example

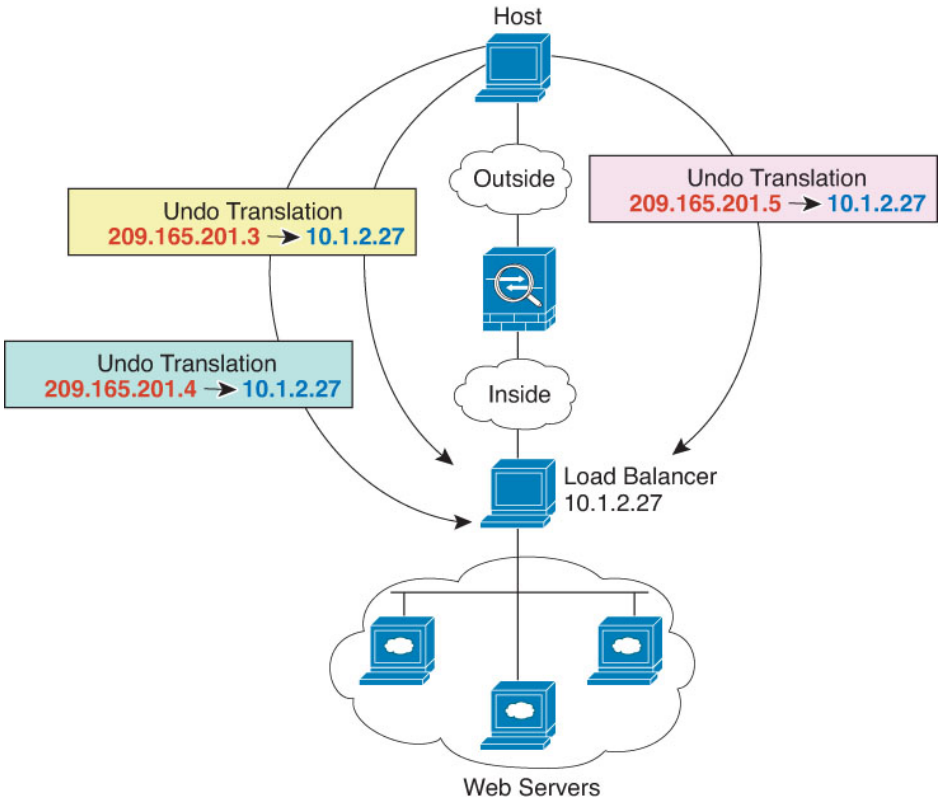

Other Mapping Scenarios (Not Recommended)

NAT has the flexibility to allow any kind of static mapping scenario: one-to-one, one-to-many, but also few-to-many, many-to-few, and many-to-one mappings. We recommend using only one-to-one or one-to-many mappings. These other mapping options might result in unintended consequences.

Functionally, few-to-many is the same as one-to-many; but because the configuration is more complicated and the actual mappings may not be obvious at a glance, we recommend creating a one-to-many configuration for each real address that requires it. For example, for a few-to-many scenario, the few real addresses are mapped to the many mapped addresses in order (A to 1, B to 2, C to 3). When all real addresses are mapped, the next mapped address is mapped to the first real address, and so on until all mapped addresses are mapped (A to 4, B to 5, C to 6). This results in multiple mapped addresses for each real address. Just like a one-to-many configuration, only the first mappings are bidirectional; subsequent mappings allow traffic to be initiated *to* the real host, but all traffic *from* the real host uses only the first mapped address for the source.

The following figure shows a typical few-to-many static NAT scenario.

Figure 401: Few-to-Many Static NAT

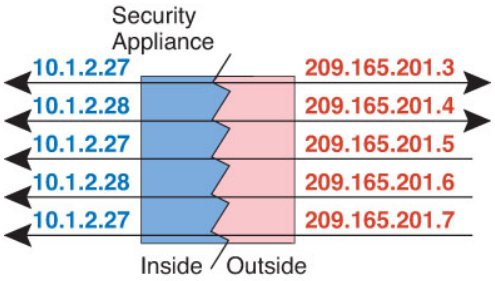

For a many-to-few or many-to-one configuration, where you have more real addresses than mapped addresses, you run out of mapped addresses before you run out of real addresses. Only the mappings between the lowest real IP addresses and the mapped pool result in bidirectional initiation. The remaining higher real addresses can initiate traffic, but traffic cannot be initiated to them (returning traffic for a connection is directed to the correct real address because of the unique 5-tuple (source IP, destination IP, source port, destination port, protocol) for the connection).

Note Many-to-few or many-to-one NAT is not PAT. If two real hosts use the same source port number and go to the same outside server and the same TCP destination port, and both hosts are translated to the same IP address, then both connections will be reset because of an address conflict (the 5-tuple is not unique).

The following figure shows a typical many-to-few static NAT scenario.

Figure 402: Many-to-Few Static NAT

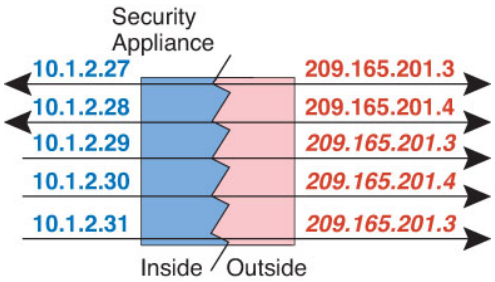

Instead of using a static rule this way, we suggest that you create a one-to-one rule for the traffic that needs bidirectional initiation, and then create a dynamic rule for the rest of your addresses.

Configure Static Auto NAT

Use static auto NAT rules to translate addresses to different IP addresses that are routable on the destination network. You can also do port translation with the static NAT rule.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Source**—This must be a network object (not a group), and it can be a host, range, or subnet.
- **Translated Source**—You have the following options to specify the translated address:
  - **Destination Interface**—To use the destination interface address, you do not need a network object. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
  - **Address**—Create a network object or group containing hosts, ranges, or subnets. A group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Auto NAT Rule**.
- **Type**—Select **Static**.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 On **Translation**, configure the following options:

- **Original Source**—The network object that contains the addresses you are translating.
- **Translated Source**—One of the following:
  - To use a set group of addresses, select **Address** and the network object or group that contains the mapped addresses. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.
  - (Static interface NAT with port translation.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
- (Optional.) **Original Port**, **Translated Port**—If you need to translate a TCP or UDP port, select the protocol in **Original Port**, and type the original and translated port numbers. For example, you can translate TCP/80 to 8080 if necessary.

Step 6 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see Rewriting DNS Queries and Responses Using NAT, on page 1117. This option is not available if you are doing port translation.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.
- **Net to Net Mapping**—For NAT 46, select this option to translate the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used. For a one-to-one translation, you must use this option.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.

Step 7 Click **Save** to add the rule.

Step 8 Click **Save** on the NAT page to save your changes.

Configure Static Manual NAT

Use static manual NAT rules when auto NAT does not meet your needs. For example, if you want to do different translations based on the destination. Static NAT translates addresses to different IP addresses that are routable on the destination network. You can also do port translation with the static NAT rule.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source**—This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source**—You have the following options to specify the translated address:
  - **Destination Interface**—To use the destination interface address, you do not need a network object. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
  - **Address**—Create a network object or group containing hosts, range, or subnets. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.

You can also create network objects or groups for the **Original Destination** and **Translated Destination** if you are configuring a static translation for those addresses in the rule. If you want to configure destination static interface NAT with port translation only, you can skip adding an object for the destination mapped addresses and specify the interface in the rule.

You can also perform port translation on the source, destination, or both. In the Object Manager, ensure that there are port objects you can use for the original and translated ports.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Manual NAT Rule**.
- **Type**—Select **Static**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- **Enable**—Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page.
- **Insert**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic

exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 (On the **Translation** page.) Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet.

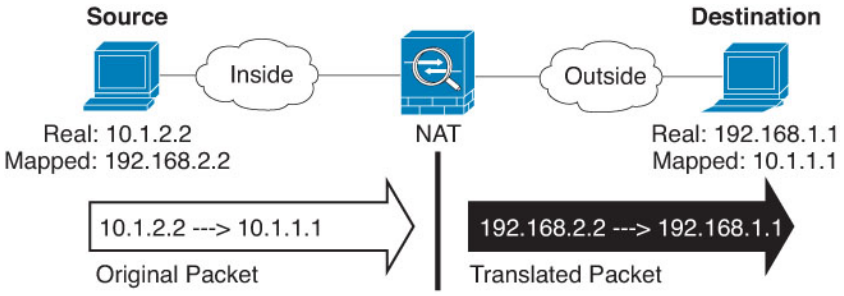

- **Original Source**—The network object or group that contains the addresses you are translating.
- **Original Destination**—(Optional.) The network object or group that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

  You can select **Source Interface IP** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source**—One of the following:
  - To use a set group of addresses, select **Address** and the network object or group that contains the mapped addresses. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.
  - (Static interface NAT with port translation.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
- **Translated Destination**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the source or destination service ports for service translation.

If you are configuring static NAT with port translation, you can translate ports for the source, destination, or both. For example, you can translate between TCP/80 and TCP/8080.

NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

- **Original Source Port**, **Translated Source Port**—Defines a port translation for the source address.
- **Original Destination Port**, **Translated Destination Port**—Defines a port translation for the destination address.

Step 8 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see Rewriting DNS Queries and Responses Using NAT, on page 1117. This option is not available if you are doing port translation.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.
- **Net to Net Mapping**—For NAT 46, select this option to translate the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used. For a one-to-one translation, you must use this option.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.
- **Unidirectional**—Select this option to prevent the destination addresses from initiating traffic to the source addresses. The unidirectional option is mainly useful for testing purposes and might not work with all protocols. For example, SIP requires protocol inspection to translate SIP headers using NAT, but this will not occur if you make the translation unidirectional.

Step 9 Click **Save** to add the rule.

Step 10 Click **Save** on the NAT page to save your changes.

Identity NAT

You might have a NAT configuration in which you need to translate an IP address to itself. For example, if you create a broad rule that applies NAT to every network, but want to exclude one network from NAT, you can create a static NAT rule to translate an address to itself.

The following figure shows a typical identity NAT scenario.

Figure 403: Identity NAT

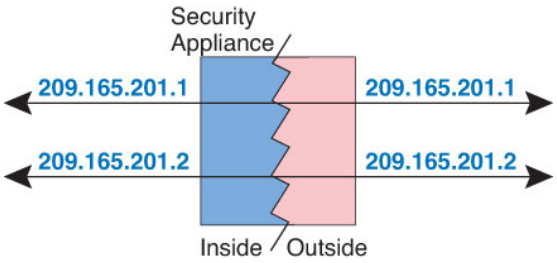

The following topics explain how to configure identity NAT.

Configure Identity Auto NAT

Use static identity auto NAT rules to prevent the translation of an address. That is, to translate the address to itself.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Alternatively, you can create the objects while defining the NAT rule. The objects must meet the following requirements:

- **Original Source**—This must be a network object (not a group), and it can be a host, range, or subnet.
- **Translated Source**—A network object or group with the exact same contents as the original source object. You can use the same object.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Auto NAT Rule**.
- **Type**—Select **Static**.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 On **Translation**, configure the following options:

- **Original Source**—The network object that contains the addresses you are translating.
- **Translated Source**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.

Do not configure the **Original Port** and **Translated Port** options for identity NAT.

Step 6 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Do not configure this option for identity NAT.
- **IPv6**—Do not configure this option for identity NAT.
- **Net to Net Mapping**—Do not configure this option for identity NAT.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.
- **Perform Route Lookup for Destination Interface**— If you select source and destination interfaces when selecting the same object for original and translated source address, you can select this option to have the system determine the destination interface based on the routing table rather than using the destination interface configured in the NAT rule.

Step 7 Click **Save** to add the rule.

Step 8 Click **Save** on the NAT page to save your changes.

Configure Identity Manual NAT

Use static identity manual NAT rules when auto NAT does not meet your needs. For example, if you want to do different translations based on the destination. Use static identity NAT rules to prevent the translation of an address. That is, to translate the address to itself.

Before you begin

Select **Objects** > **Object Management** and create the network objects or groups needed in the rule. Groups cannot contain both IPv4 and IPv6 addresses; they must contain one type only. Alternatively, you can create the objects while defining the NAT rule. The objects must also meet the following requirements:

- **Original Source**—This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can skip this step and specify **Any** in the rule.
- **Translated Source**—The same object or group as the original source. Optionally, you can select a different object that has the exact same contents.

You can also create network objects or groups for the **Original Destination** and **Translated Destination** if you are configuring a static translation for those addresses in the rule. If you want to configure destination static interface NAT with port translation only, you can skip adding an object for the destination mapped addresses and specify the interface in the rule.

You can also perform port translation on the source, destination, or both. In the Object Manager, ensure that there are port objects you can use for the original and translated ports. You can use the same object for identity NAT.

Procedure

Step 1 Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

Step 2 Do one of the following:

- Click the **Add Rule** button to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The right click menu also has options to cut, copy, paste, insert, and delete rules.

Step 3 Configure the basic rule options:

- **NAT Rule**—Select **Manual NAT Rule**.
- **Type**—Select **Static**. This setting only applies to the source address. If you define a translation for the destination address, the translation is always static.
- **Enable**—Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page.
- **Insert**—Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Step 4 On **Interface Objects**, configure the following options:

- **Source Interface Objects**, **Destination Interface Objects**—(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Step 5 Identify the original packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear in the original packet.

See the following figure for an example of the original packet vs. the translated packet where you perform identity NAT on the inside host but translate the outside host.

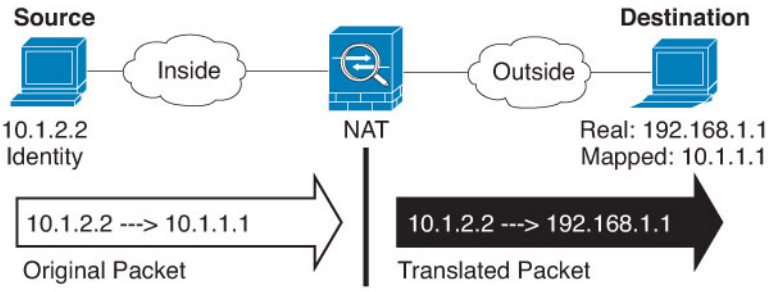

- **Original Source**—The network object or group that contains the addresses you are translating.
- **Original Destination**—(Optional.) The network object or group that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If

you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

You can select **Interface Object** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Step 6 Identify the translated packet addresses, either IPv4 or IPv6; namely, the packet addresses as they appear on the destination interface network. You can translate between IPv4 and IPv6 if desired.

- **Translated Source**—The same object or group as the original source. Optionally, you can select a different object that has the exact same contents.
- **Translated Destination**—(Optional.) The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

Step 7 (Optional.) Identify the source or destination service ports for service translation.

If you are configuring static NAT with port translation, you can translate ports for the source, destination, or both. For example, you can translate between TCP/80 and TCP/8080.

NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

- **Original Source Port**, **Translated Source Port**—Defines a port translation for the source address.
- **Original Destination Port**, **Translated Destination Port**—Defines a port translation for the destination address.

Step 8 (Optional.) On **Advanced**, select the desired options:

- **Translate DNS replies that match this rule**—Do not configure this option for identity NAT.
- **IPv6**—Whether to use the IPv6 address of the destination interface for interface PAT.
- **Do not proxy ARP on Destination Interface**—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.
- **Perform Route Lookup for Destination Interface**— If you select source and destination interfaces when selecting the same object for original and translated source address, you can select this option to have the system determine the destination interface based on the routing table rather than using the destination interface configured in the NAT rule.
- **Unidirectional**—Select this option to prevent the destination addresses from initiating traffic to the source addresses. The unidirectional option is mainly useful for testing purposes and might not work with all protocols. For example, SIP requires protocol inspection to translate SIP headers using NAT, but this will not occur if you make the translation unidirectional.

Step 9 Click **Save** to add the rule.

Step 10 Click **Save** on the NAT page to save your changes.

NAT Rule Properties for Firewall Threat Defense

Use Network Address Translation (NAT) rules to translate IP addresses to other IP addresses. You would typically use NAT rules to convert private addresses to publicly routable addresses. The translation can be from one address to another, or you can use Port Address Translation (PAT) to translate many addresses to one or a few addresses, using port numbers to distinguish among the source addresses.

NAT rules include the following basic properties. The properties are the same for auto NAT and manual NAT rules except where indicated.

NAT Type

Whether you want to configure a **Manual NAT Rule** or an **Auto NAT Rule**. Auto NAT translates the source address only, and you cannot make different translations based on the destination address. Because auto NAT is more simple to configure, use it unless you need the added features of manual NAT. For more information on the differences, see Auto NAT and Manual NAT, on page 1019.

Type

Whether the translation rule is **Dynamic** or **Static**. Dynamic translation automatically chooses the mapped address from a pool of addresses, or an address/port combination when implementing PAT. Use static translation if you want to precisely define the mapped address/port.

Enable (Manual NAT only.)

Whether you want the rule to be active. You can later activate or deactivate the rule using the right-click menu on the rules page. You cannot disable auto NAT rules.

Insert (Manual NAT only.)

Where you want to add the rule. You can insert it in a category (before or after auto NAT rules), or above or below the rule number you specify.

Description (Optional. Manual NAT only.)

A description of the purpose of the rule.

The following topics describe the tabs for the NAT rules properties.

Interface Objects NAT Properties

Interface objects (security zones or interface groups) define the interfaces to which a NAT rule applies. In routed mode, you can use the default "any" for both source and destination to apply to all interfaces of all assigned devices. However, you typically want to select specific source and destination interfaces.

Notes

- The concept of “any” interface does not apply to bridge group member interfaces. When you specify “any” interface, all bridge group member interfaces are excluded. Thus, to apply NAT to bridge group members, you must specify the member interface. You cannot configure NAT for the Bridge Virtual Interface (BVI) itself, you can configure NAT for member interfaces only.

  If you select interface objects, a NAT rule will be configured on an assigned device only if the device has interfaces included in all selected objects. For example, if you select both source and destination security zones, both zones must contain one or more interface for a given device.

- If more than one interface in an interface object exists on a given device, identical rules are created for each interface. This can become an issue for static NAT rules that include destination translation. Because NAT rules are applied based on first hit rule, only the rule created for the first interface configured for

the object matches traffic. When configuring static NAT with destination translation, use interface objects that include at most one interface per device assigned to the NAT policy to ensure you are getting the desired results.

Source Interface Objects, Destination Interface Objects

(Required for bridge group member interfaces.) The interface objects (security zones or interface groups) that identify the interfaces where this NAT rule applies. **Source** is the object containing the real interface, the one through which the traffic enters the device. **Destination** is the object containing the mapped interface, the one through which traffic exits the device. By default, the rule applies to all interfaces (**Any**) except for bridge group member interfaces.

Translation Properties for Auto NAT

Use the **Translation** options to define the source addresses and the mapped translated addresses. The following properties apply to auto NAT only.

Original Source (Always required.)

The network object that contains the addresses you are translating. This must be a network object (not a group), and it can be a host, range, or subnet.

You cannot create auto NAT rules for the system-defined any-ipv4 or any-ipv6 objects.

Translated Source (Usually required.)

The mapped addresses, the ones to which you are translating. What you select here depends on the type of translation rule you are defining.

- **Dynamic NAT**—The network object or group that contains the mapped addresses. This can be a network object or group, but it cannot include a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. If a group contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback.
- **Dynamic PAT**—One of the following:
  - (Interface PAT.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. Do not configure a PAT pool.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose. Do not configure a PAT pool.
  - To use a PAT pool, leave **Translated Source** empty. Select the PAT pool object on **PAT Pool**.
- **Static NAT**—One of the following:
  - To use a set group of addresses, select **Address** and the network object or group that contains the mapped addresses. The object or group can contain hosts, ranges, or subnets. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.
  - (Static interface NAT with port translation.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on the **Advanced** tab. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.

- **Identity NAT**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.

Original Port, Translated Port (Static NAT only.)

If you need to translate a TCP or UDP port, select the protocol in **Original Port**, and type the original and translated port numbers. For example, you can translate TCP/80 to 8080 if necessary. Do not configure these options for identity NAT.

Translation Properties for Manual NAT

Use the **Translation** options to define the source addresses and the mapped translated addresses. The following properties apply to manual NAT only. All are optional except as indicated.

Original Source (Always required.)

The network object or group that contains the addresses you are translating. This can be a network object or group, and it can contain a host, range, or subnet. If you want to translate all original source traffic, you can specify **Any** in the rule.

Translated Source (Usually required.)

The mapped addresses, the ones to which you are translating. What you select here depends on the type of translation rule you are defining.

- **Dynamic NAT**—The network object or group that contains the mapped addresses. This can be a network object or group, but it cannot include a subnet. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only. If a group contains both ranges and host IP addresses, then the ranges are used for dynamic NAT, and then the host IP addresses are used as a PAT fallback.
- **Dynamic PAT**—One of the following:
  - (Interface PAT.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on **Advanced**. Do not configure a PAT pool.
  - To use a single address other than the destination interface address, select the host network object you created for this purpose. Do not configure a PAT pool.
  - To use a PAT pool, leave **Translated Source** empty. Select the PAT pool object on **PAT Pool**.
- **Static NAT**—One of the following:
  - To use a set group of addresses, select **Address** and the network object or group that contains the mapped addresses. The object or group can contain hosts, ranges, or subnets. Typically, you configure the same number of mapped addresses as real addresses for a one-to-one mapping. You can, however, have a mismatched number of addresses.
  - (Static interface NAT with port translation.) To use the address of the destination interface, select **Destination Interface IP**. You must also select a specific destination interface object. To use the IPv6 address of the interface, you must also select the **IPv6** option on the **Advanced** tab. This configures static interface NAT with port translation: the source address/port is translated to the interface's address and the same port number.
- **Identity NAT**—The same object as the original source. Optionally, you can select a different object that has the exact same contents.

Original Destination

The network object or group that contains the addresses of the destinations. If you leave this blank, the source address translation applies regardless of destination. If you do specify the destination address, you can configure a static translation for that address or just use identity NAT for it.

You can select **Source Interface IP** to base the original destination on the source interface (which cannot be Any). If you select this option, you must also select a translated destination object. To implement a static interface NAT with port translation for the destination addresses, select this option and also select the appropriate port objects for the destination ports.

Translated Destination

The network object or group that contains the destination addresses used in the translated packet. If you selected an object for **Original Destination**, you can set up identity NAT (that is, no translation) by selecting the same object.

You can use a network object that specifies a fully-qualified domain name as the translated destination; for more information, see FQDN Destination Guidelines, on page 1028.

Original Source Port, Translated Source Port, Original Destination Port, Translated Destination Port

The port objects that define the source and destination services for the original and translated packets. You can translate the ports, or select the same object to make the rule sensitive to the service without translating the ports. Keep the following rules in mind when configuring services:

- (Dynamic NAT or PAT.) You cannot do translation on the **Original Source Port** and **Translated Source Port**. You can do translation on the destination port only.
- NAT only supports TCP or UDP. When translating a port, be sure the protocols in the real and mapped service objects are identical (both TCP or both UDP). For identity NAT, you can use the same service object for both the real and mapped ports.

PAT Pool NAT Properties

When you configure dynamic NAT, you can define a pool of addresses to use for Port Address Translation using the properties on the **PAT Pool** tab.

Enable PAT Pool

Select this option to configure a pool of addresses for PAT.

PAT

The addresses to use for the PAT pool, one of the following:

- **Address**—The object that defines the PAT pool addresses, either a network object that includes a range, or a network object group that contains hosts, ranges, or both. You cannot include subnets. The group cannot contain both IPv4 and IPv6 addresses; it must contain one type only.
- **Destination Interface IP**—Indicates that you want to use the destination interface as the PAT address. For this option, you must select a specific **Destination Interface Object**; you cannot use **Any** as the destination interface. This is another way to implement interface PAT.

Round Robin

To assign addresses/ports in a round-robin fashion. By default without round robin, all ports for a PAT address will be allocated before the next PAT address is used. The round-robin method assigns one

address/port from each PAT address in the pool before returning to use the first address again, and then the second address, and so on.

Extended PAT Table

To use extended PAT. Extended PAT uses 65535 ports per *service*, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80. You cannot use this option with interface PAT or interface PAT fallback.

Flat Port Range; Include Reserved Ports

To use the 1024 to 65535 port range as a single flat range when allocating TCP/UDP ports. (Pre-6.7) When choosing the mapped port number for a translation, PAT uses the real source port number if it is available. However, without this option, if the real port is *not* available, by default the mapped ports are chosen from the same range of ports as the real port number: 1 to 511, 512 to 1023, and 1024 to 65535. To avoid running out of ports at the low ranges, configure this setting. To use the entire range of 1 to 65535, also check the **Include Reserved Ports** option. For the Firewall Threat Defense devices running version 6.7 or higher, the flat port range is always configured, whether you select the option or not. You can still select the **Include Reserved Ports** option for these systems, and that setting is honored.

Block Allocation

To enable port block allocation. For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time. If you allocate a block of ports, subsequent connections from the host use new randomly selected ports within the block. If necessary, additional blocks are allocated if the host has active connections for all ports in the original block. Port blocks are allocated in the 1024-65535 range only. Port block allocation is compatible with round robin, but you cannot use it with the extended PAT or flat port range options. You also cannot use interface PAT fallback.

Advanced NAT Properties

When you configure NAT, you can configure properties that provide specialized services in the **Advanced** options. All of these properties are optional: configure them only if you need the service.

Translate DNS replies that match this rule

Whether to translate the IP address in DNS replies. For DNS replies traversing from a mapped interface to a real interface, the Address (the IPv4 A or IPv6 AAAA) record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the record is rewritten from the real value to the mapped value. This option is used in specific circumstances, and is sometimes needed for NAT64/46 translation, where the rewrite also converts between A and AAAA records. For more information, see Rewriting DNS Queries and Responses Using NAT, on page 1117. This option is not available if you are doing port translation in a static NAT rule.

Fallthrough to Interface PAT (Destination Interface) (Dynamic NAT only.)

Whether to use the IP address of the destination interface as a backup method when the other mapped addresses are already allocated (interface PAT fallback). This option is available only if you select a destination interface that is not a member of a bridge group. To use the IPv6 address of the interface, also check the **IPv6** option. You cannot select this option if you already configured interface PAT as the translated address. You also cannot select the option if you configure a PAT pool.

IPv6

Whether to use the IPv6 address of the destination interface for interface PAT.

Net to Net Mapping (Static NAT only.)

For NAT 46, select this option to translate the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used. For a one-to-one translation, you must use this option.

Do not proxy ARP on Destination Interface (Static NAT only.)

Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router. Normally for identity NAT, proxy ARP is not required, and in some cases can cause connectivity issues.

Perform Route Lookup for Destination Interface (Static Identity NAT only. Routed mode only.)

If you select source and destination interfaces when selecting the same object for original and translated source address, you can select this option to have the system determine the destination interface based on the routing table rather than using the destination interface configured in the NAT rule.

Unidirectional (Manual NAT only, static NAT only.)

Select this option to prevent the destination addresses from initiating traffic to the source addresses. The unidirectional option is mainly useful for testing purposes and might not work with all protocols. For example, SIP requires protocol inspection to translate SIP headers using NAT, but this will not occur if you make the translation unidirectional.

Translating IPv6 Networks

In cases where you need to pass traffic between IPv6-only and IPv4-only networks, you need to use NAT to convert between the address types. Even with two IPv6 networks, you might want to hide internal addresses from the outside network.

You can use the following translation types with IPv6 networks:

- NAT64, NAT46—Translates IPv6 packets into IPv4 and vice versa. You need to define two policies, one for the IPv6 to IPv4 translation, and one for the IPv4 to IPv6 translation. Although you can accomplish this with a single manual NAT rule, if the DNS server is on the external network, you probably need to rewrite the DNS response. Because you cannot enable DNS rewrite on a manual NAT rule when you specify a destination, creating two auto NAT rules is the better solution.

Note NAT46 supports static mappings only.

- NAT66—Translates IPv6 packets to a different IPv6 address. We recommend using static NAT. Although you can use dynamic NAT or PAT, IPv6 addresses are in such large supply, you do not have to use dynamic NAT.

Note NAT64 and NAT 46 are possible on standard routed interfaces only. NAT66 is possible on both routed and bridge group member interfaces.

NAT64/46: Translating IPv6 Addresses to IPv4

When traffic goes from an IPv6 network to an IPv4-only network, you need to convert the IPv6 address to IPv4, and return traffic from IPv4 to IPv6. You need to define two address pools, an IPv4 address pool to bind IPv6 addresses in the IPv4 network, and an IPv6 address pool to bind IPv4 addresses in the IPv6 network.

- The IPv4 address pool for the NAT64 rule is normally small and typically might not have enough addresses to map one-to-one with the IPv6 client addresses. Dynamic PAT might more easily meet the possible large number of IPv6 client addresses compared to dynamic or static NAT.
- The IPv6 address pool for the NAT46 rule can be equal to or larger than the number of IPv4 addresses to be mapped. This allows each IPv4 address to be mapped to a different IPv6 address. NAT46 supports static mappings only, so you cannot use dynamic PAT.

You need to define two policies, one for the source IPv6 network, and one for the destination IPv4 network. Although you can accomplish this with a single manual NAT rule, if the DNS server is on the external network, you probably need to rewrite the DNS response. Because you cannot enable DNS rewrite on a manual NAT rule when you specify a destination, creating two auto NAT rules is the better solution.

NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet

Following is a straight-forward example where you have an inside IPv6-only network, and you want to convert to IPv4 for traffic sent to the Internet. This example assumes you do not need DNS translation, so you can perform both the NAT64 and NAT46 translations in a single manual NAT rule.

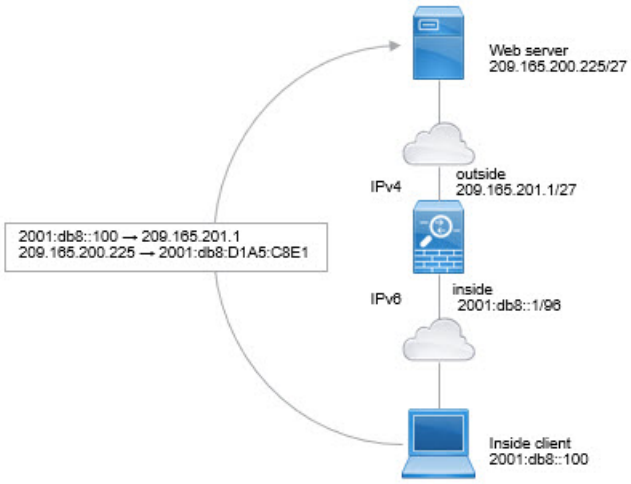

In this example, you translate the inside IPv6 network to IPv4 using dynamic interface PAT with the IP address of the outside interface. Outside IPv4 traffic is statically translated to addresses on the 2001:db8::/96 network, allowing transmission on the inside network.

Procedure

Step 1 Create the network object that defines the inside IPv6 network.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the inside IPv6 network.

Name the network object (for example, inside_v6) and enter the network address, 2001:db8::/96.

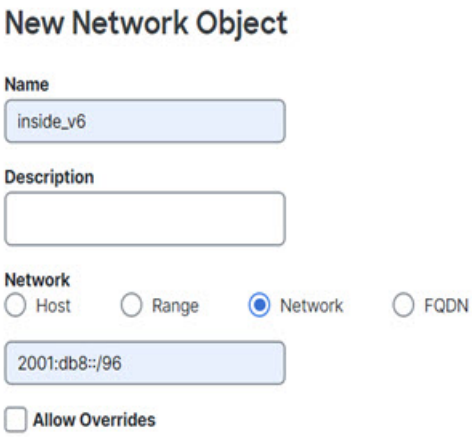

d) Click **Save**.

Step 2 Create the manual NAT rule to translate the IPv6 network to IPv4 and back again.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = inside_v6 network object.
- **Translated Source** = **Destination Interface IP**.
- **Original Destination** = inside_v6 network object.
- **Translated Destination** = any-ipv4 network object.

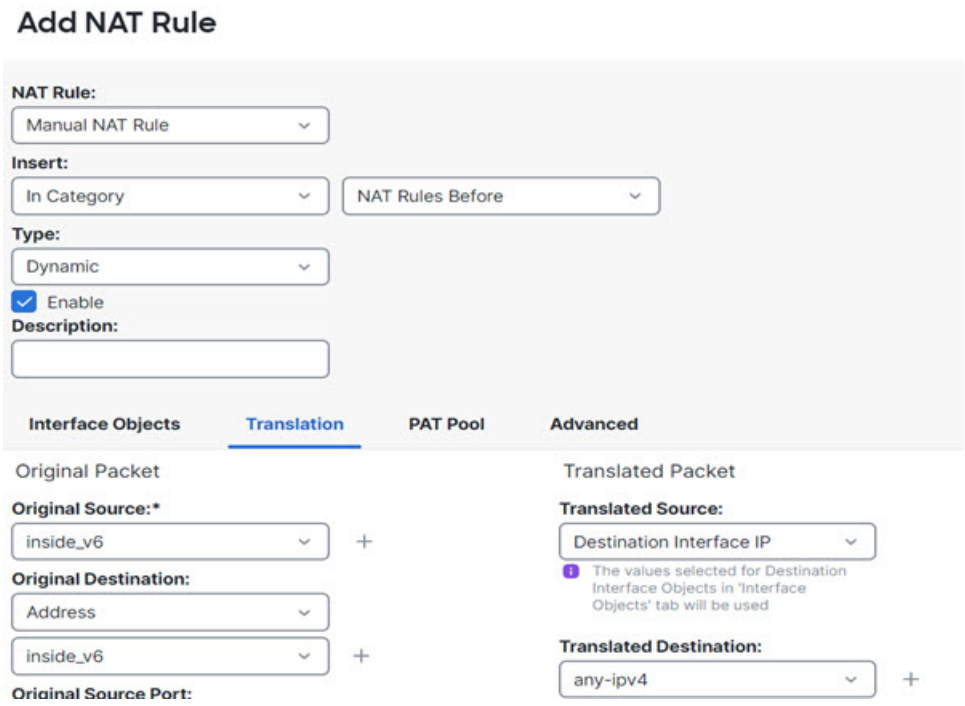

f) Click **OK**.

With this rule, any traffic from the 2001:db8::/96 subnet on the inside interface going to the outside interface gets a NAT64 PAT translation using the IPv4 address of the outside interface. Conversely, any IPv4 address on the outside network coming to the inside interface is translated to an address on the 2001:db8::/96 network using the embedded IPv4 address method.

g) Click **Save** on the NAT rules page.

NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet and DNS Translation

Following is a typical example where you have an inside IPv6-only network, but there are some IPv4-only services on the outside Internet that internal users need.

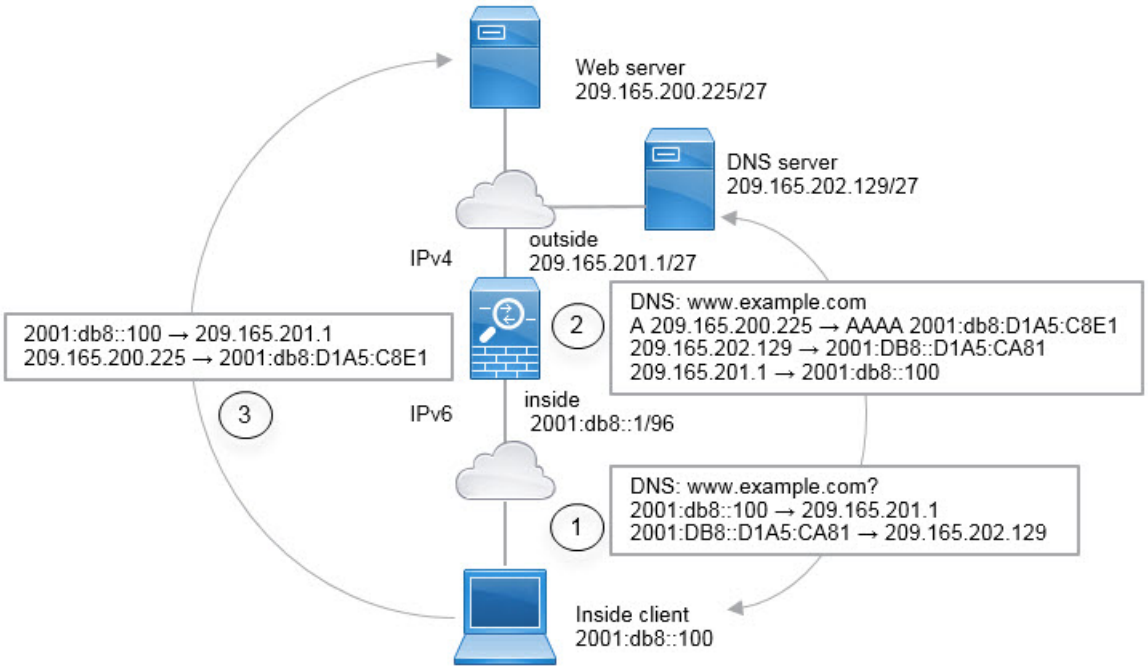

In this example, you translate the inside IPv6 network to IPv4 using dynamic interface PAT with the IP address of the outside interface. Outside IPv4 traffic is statically translated to addresses on the 2001:db8::/96 network, allowing transmission on the inside network. You enable DNS rewrite on the NAT46 rule, so that replies from the external DNS server can be converted from A (IPv4) to AAAA (IPv6) records, and the addresses converted from IPv4 to IPv6.

Following is a typical sequence for a web request where a client at 2001:DB8::100 on the internal IPv6 network tries to open www.example.com.

1. The client's computer sends a DNS request to the DNS server at 2001:DB8::D1A5:CA81. The NAT rules make the following translations to the source and destination in the DNS request:
   - 2001:DB8::100 to a unique port on 209.165.201.1 (The NAT64 interface PAT rule.)
   - 2001:DB8::D1A5:CA81 to 209.165.202.129 (The NAT46 rule. D1A5:CA81 is the IPv6 equivalent of 209.165.202.129.)
2. The DNS server responds with an A record indicating that www.example.com is at 209.165.200.225. The NAT46 rule, with DNS rewrite enabled, converts the A record to the IPv6-equivalent AAAA record, and translates 209.165.200.225 to 2001:db8:D1A5:C8E1in the AAAA record. In addition, the source and destination addresses in the DNS response are untranslated:
   - 209.165.202.129 to 2001:DB8::D1A5:CA81
   - 209.165.201.1 to 2001:db8::100
3. The IPv6 client now has the IP address of the web server, and makes an HTTP request to www.example.com at 2001:db8:D1A5:C8E1. (D1A5:C8E1 is the IPv6 equivalent of 209.165.200.225.) The source and destination of the HTTP request are translated:
   - 2001:DB8::100 to a unique port on 209.156.101.54 (The NAT64 interface PAT rule.)
   - 2001:db8:D1A5:C8E1 to 209.165.200.225 (The NAT46 rule.)

The following procedure explains how to configure this example.

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

Step 1 Create the network objects that define the inside IPv6 and outside IPv4 networks.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the inside IPv6 network.

Name the network object (for example, inside_v6) and enter the network address, 2001:db8::/96.

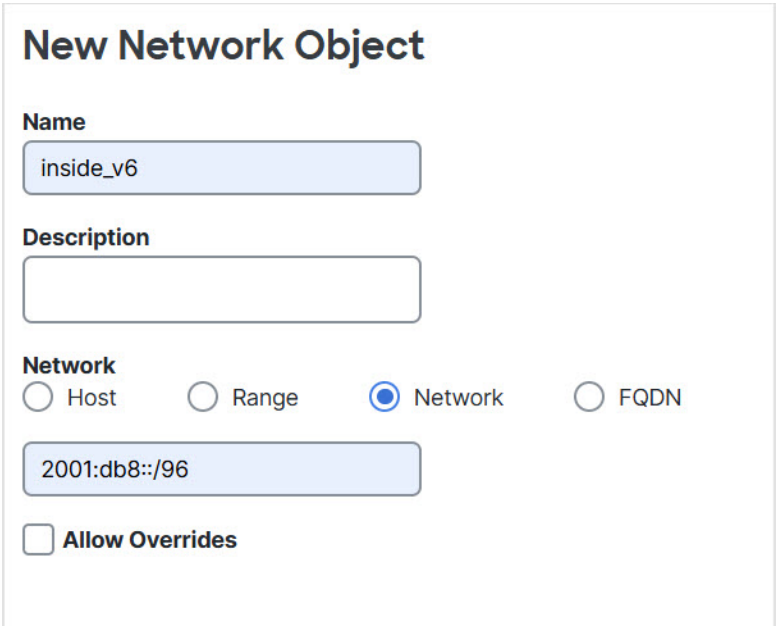

d) Click **Save**.

e) Click **Add Network** > **Add Object** and define the outside IPv4 network.

Name the network object (for example, outside_v4_any) and enter the network address 0.0.0.0/0.

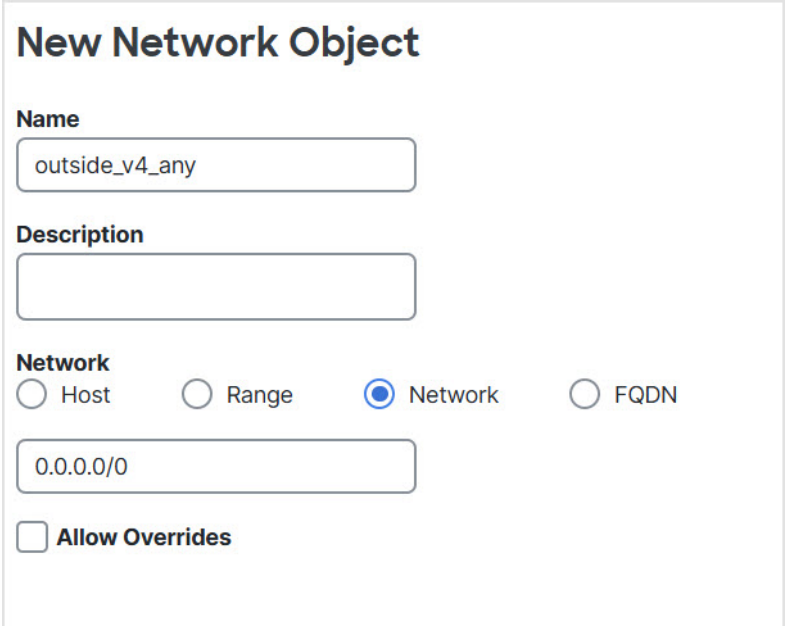

f) Click **Save**.

Step 2 Configure the NAT64 dynamic PAT rule for the inside IPv6 network.

Step 3 Configure the static NAT46 rule for the outside IPv4 network.

a) Click **Add Rule**.

b) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.
- **Destination Interface Objects** = inside.

d) On **Translation**, configure the following:

- **Original Source** = outside_v4_any network object.
- **Translated Source** > **Address** = inside_v6 network object.

e) On **Advanced**, select **Translate DNS replies that match this rule**.

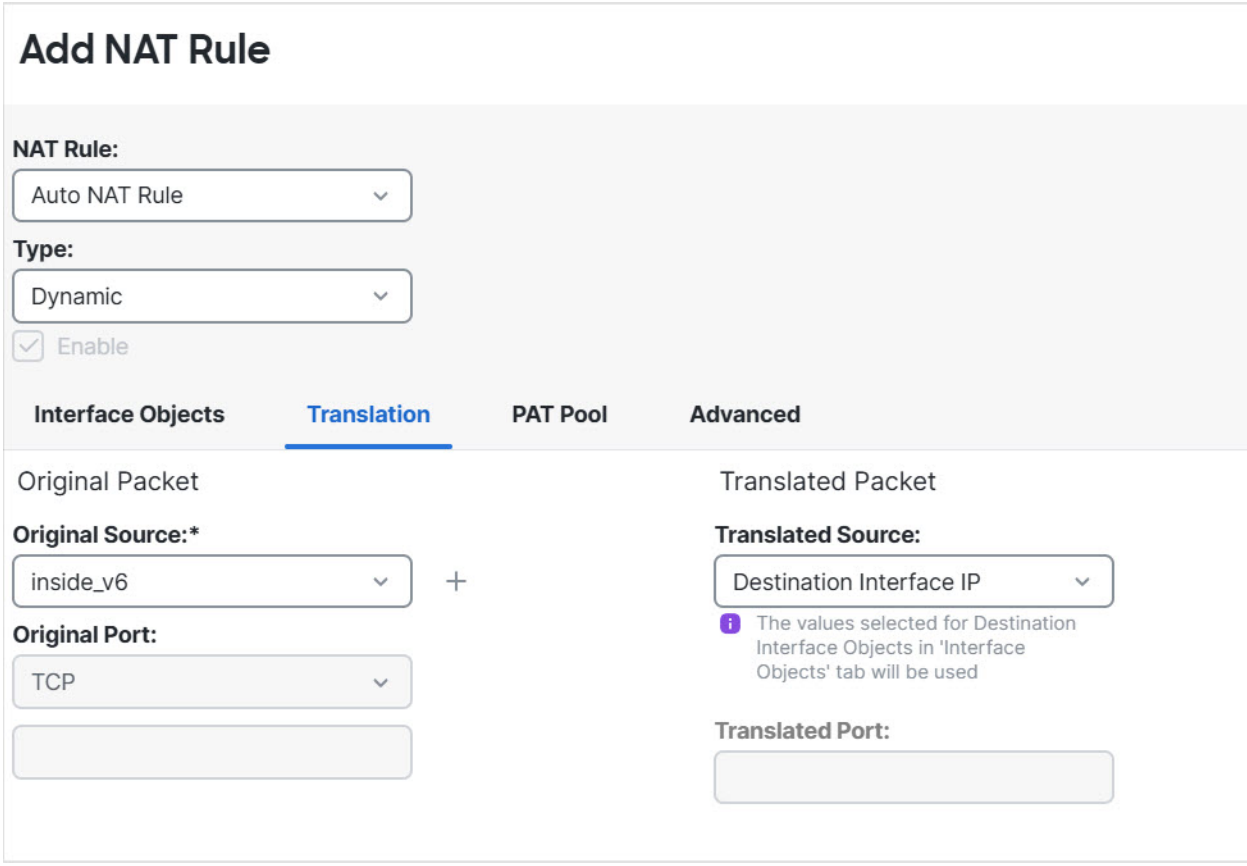

f) Click **OK**.

With this rule, any IPv4 address on the outside network coming to the inside interface is translated to an address on the 2001:db8::/96 network using the embedded IPv4 address method. In addition, DNS responses are converted from A (IPv4) to AAAA (IPv6) records, and the addresses converted from IPv4 to IPv6.

NAT66: Translating IPv6 Addresses to Different IPv6 Addresses

When going from an IPv6 network to another IPv6 network, you can translate the addresses to different IPv6 addresses on the outside network. We recommend using static NAT. Although you can use dynamic NAT or PAT, IPv6 addresses are in such large supply, you do not have to use dynamic NAT.

Because you are not translating between different address types, you need a single rule for NAT66 translations. You can easily model these rules using auto NAT. However, if you do not want to allow returning traffic, you can make the static NAT rule unidirectional using manual NAT only.

NAT66 Example, Static Translation between Networks

You can configure a static translation between IPv6 address pools using auto NAT. The following example explains how to convert inside addresses on the 2001:db8:122:2091::/96 network to outside addresses on the 2001:db8:122:2999::/96 network.

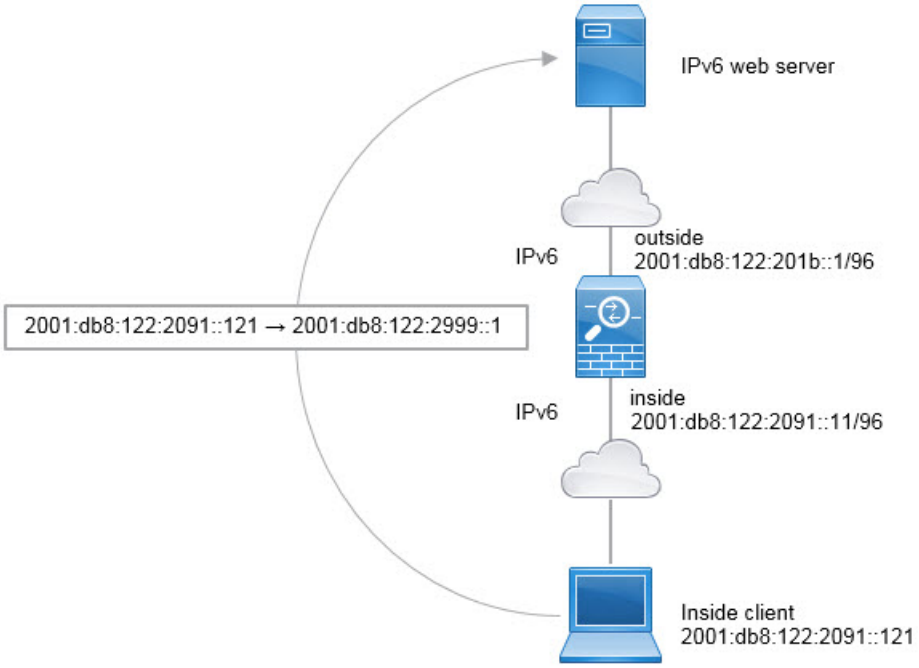

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

Step 1 Create the network objects that define the inside IPv6 and outside IPv6 NAT networks.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the inside IPv6 network.

Name the network object (for example, inside_v6) and enter the network address, 2001:db8:122:2091::/96.

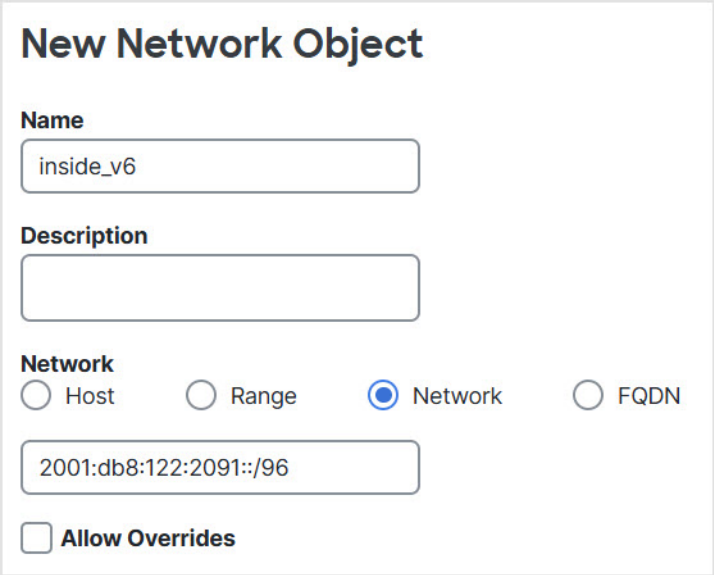

d) Click **Save**.

e) Click **Add Network** > **Add Object** and define the outside IPv6 NAT network.

Name the network object (for example, outside_nat_v6) and enter the network address 2001:db8:122:2999::/96.

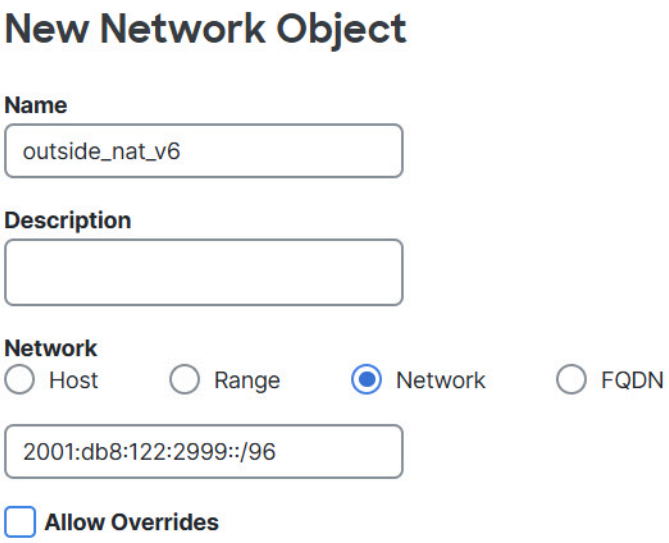

f) Click **Save**.

Step 2 Configure the static NAT rule for the inside IPv6 network.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.

- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = inside_v6 network object.
- **Translated Source** > **Address** = outside_nat_v6 network object.

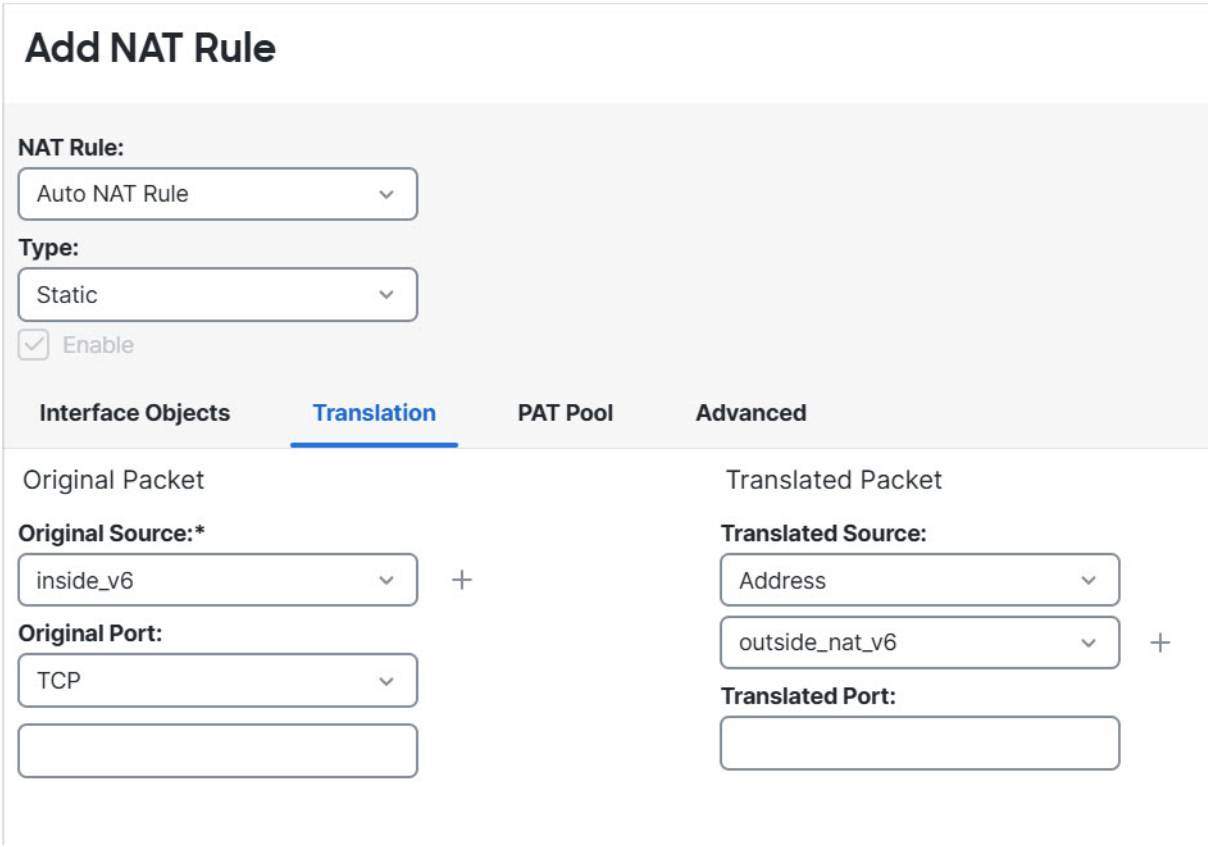

f) Click **OK**.

With this rule, any traffic from the 2001:db8:122:2091::/96 subnet on the inside interface going to the outside interface gets a static NAT66 translation to an address on the 2001:db8:122:2999::/96 network.

NAT66 Example, Simple IPv6 Interface PAT

A simple approach for implementing NAT66 is to dynamically assign internal addresses to different ports on the outside interface IPv6 address.

When you configure an interface PAT rule for NAT66, all the global addresses that are configured on that interface are used for PAT mapping. Link-local or site-local addresses for the interface are not used for PAT.

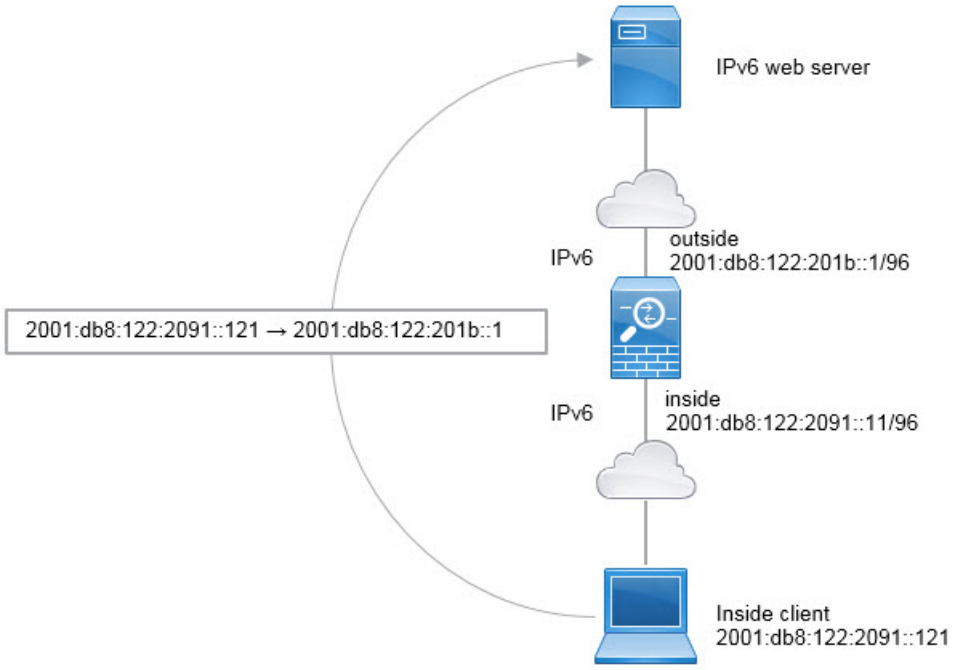

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

Step 1 Create the network object that defines the inside IPv6 network.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the inside IPv6 network.

Name the network object (for example, inside_v6) and enter the network address, 2001:db8:122:2091::/96.

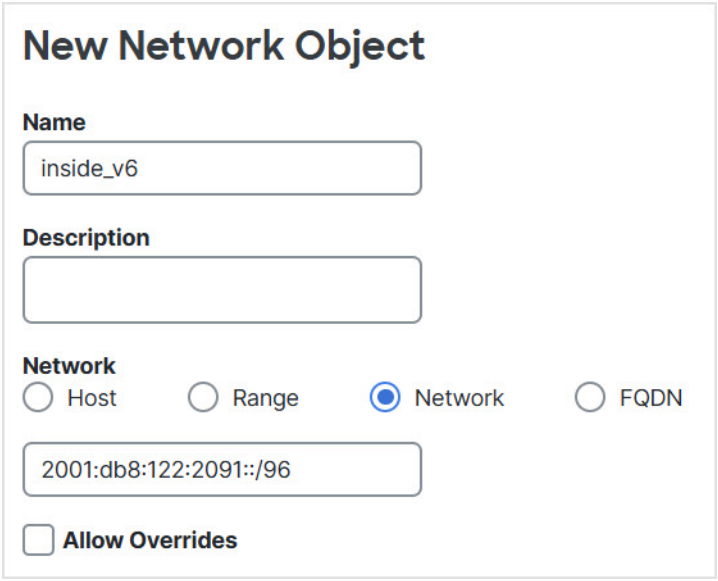

d) Click **Save**.

Step 2 Configure the dynamic PAT rule for the inside IPv6 network.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Dynamic.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = inside_v6 network object.
- **Translated Source** = **Destination Interface IP**.

f) On **Advanced**, select **IPv6**, which indicates that the IPv6 address of the destination interface should be used.

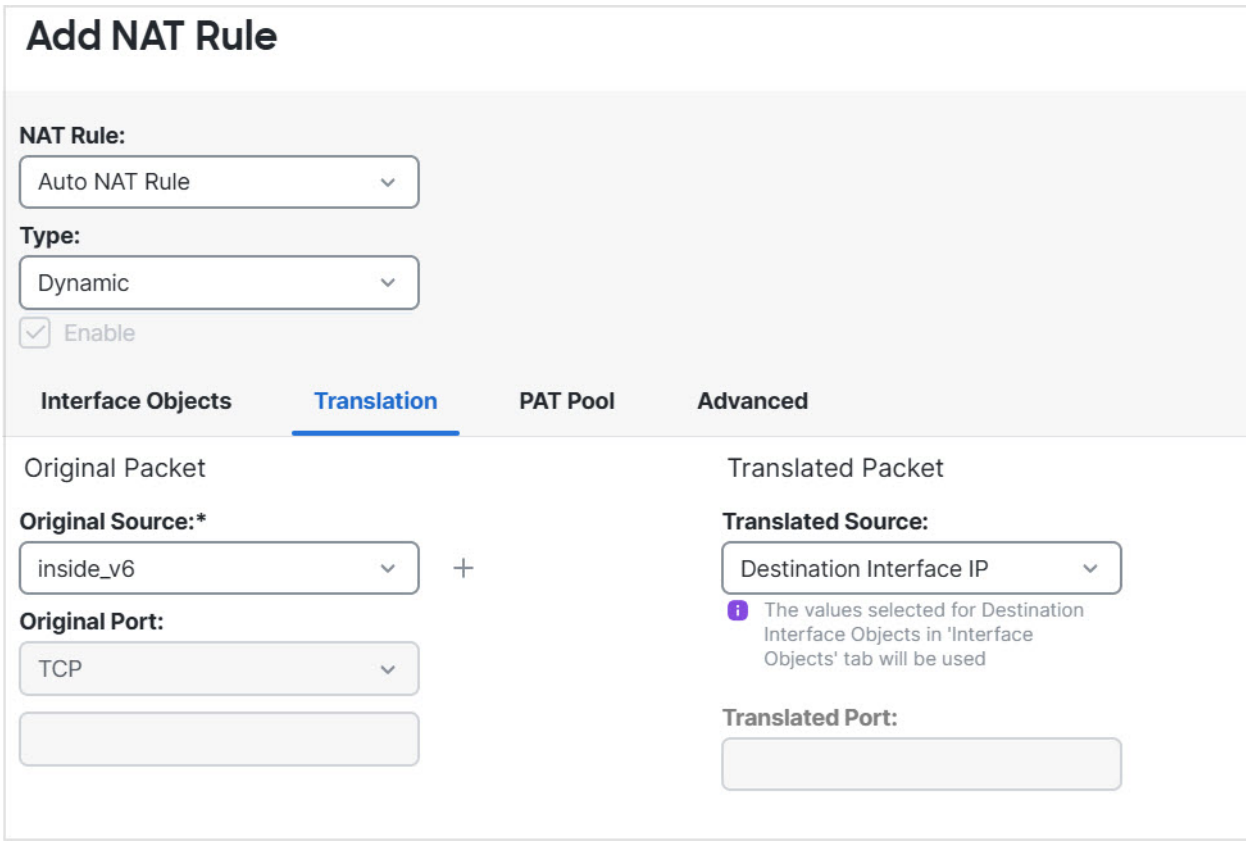

g) Click **OK**.

With this rule, any traffic from the 2001:db8:122:2091::/96 subnet on the inside interface going to the outside interface gets a NAT66 PAT translation to one of the IPv6 global addresses configured for the outside interface.

Monitoring NAT

To monitor and troubleshoot NAT connections, log into the device CLI and use the following commands.

- **show nat** displays the NAT rules and per-rule hit counts. There are additional keywords to show other aspects of NAT.
- **show xlate** displays the actual NAT translations that are currently active.
- **clear xlate** lets you remove an active NAT translation. You might need to remove active translations if you alter NAT rules, because existing connections continue to use the old translation slot until the connection ends. Clearing a translation allows the system to build a new translation for a client on the client's next connection attempt based on your new rules.

Examples for NAT

The following topics provide examples for configuring NAT on Threat Defense devices.

Providing Access to an Inside Web Server (Static Auto NAT)

The following example performs static NAT for an inside web server. The real address is on a private network, so a public address is required. Static NAT is necessary so hosts can initiate traffic to the web server at a fixed address.

Figure 404: Static NAT for an Inside Web Server

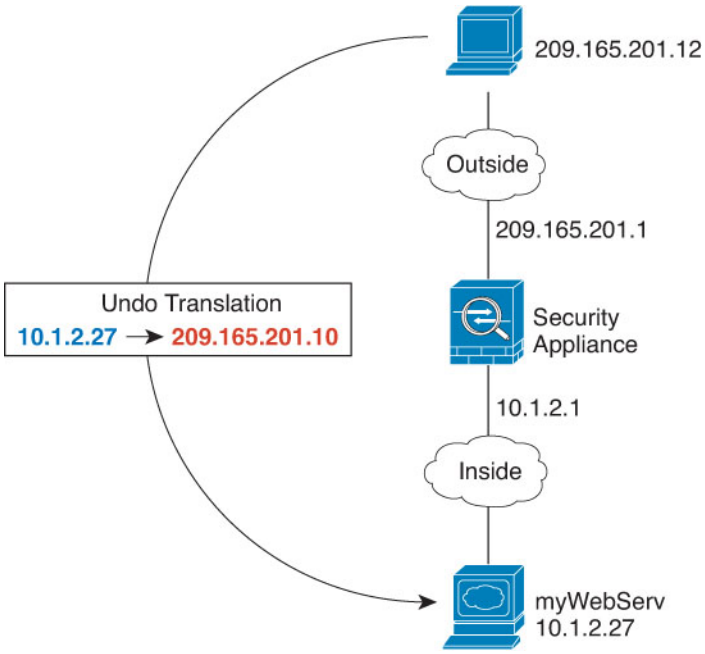

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the web server. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

Step 1 Create the network objects that define the server's private and public host addresses.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the web server's private address.

Name the network object (for example, WebServerPrivate) and enter the real host IP address, 10.1.2.27.

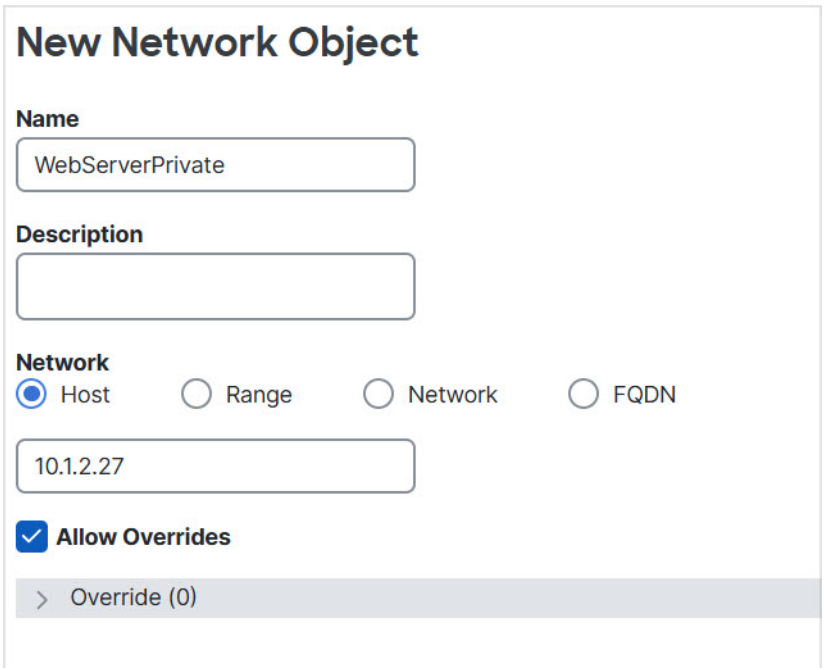

d) Click **Save**.

e) Click **Add Network** > **Add Object** and define the public address.

Name the network object (for example, WebServerPublic) and enter the host address 209.165.201.10.

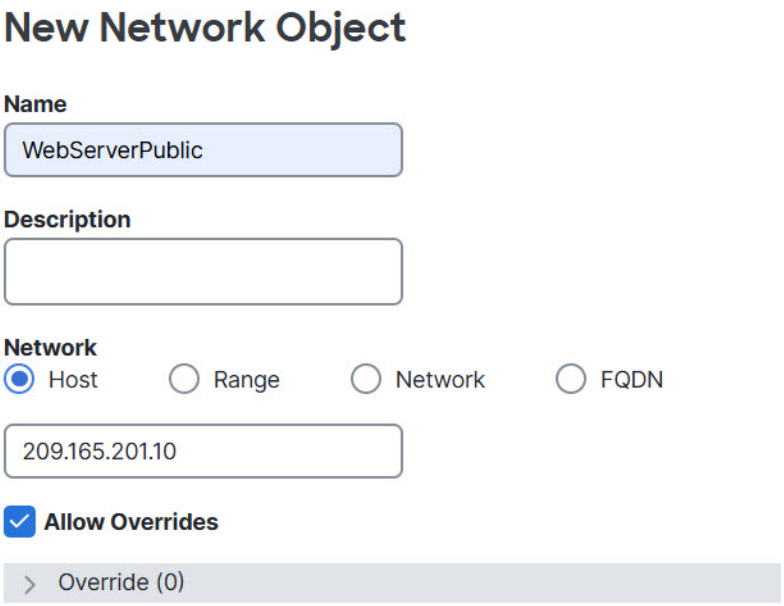

f) Click **Save**.

Step 2 Configure static NAT for the object.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.
c) Configure the following properties:
   - **NAT Rule** = Auto NAT Rule.
   - **Type** = Static.
d) On **Interface Objects**, configure the following:
   - **Source Interface Objects** = inside.
   - **Destination Interface Objects** = outside.
e) On **Translation**, configure the following:
   - **Original Source** = WebServerPrivate network object.
   - **Translated Source** > **Address**= WebServerPublic network object.

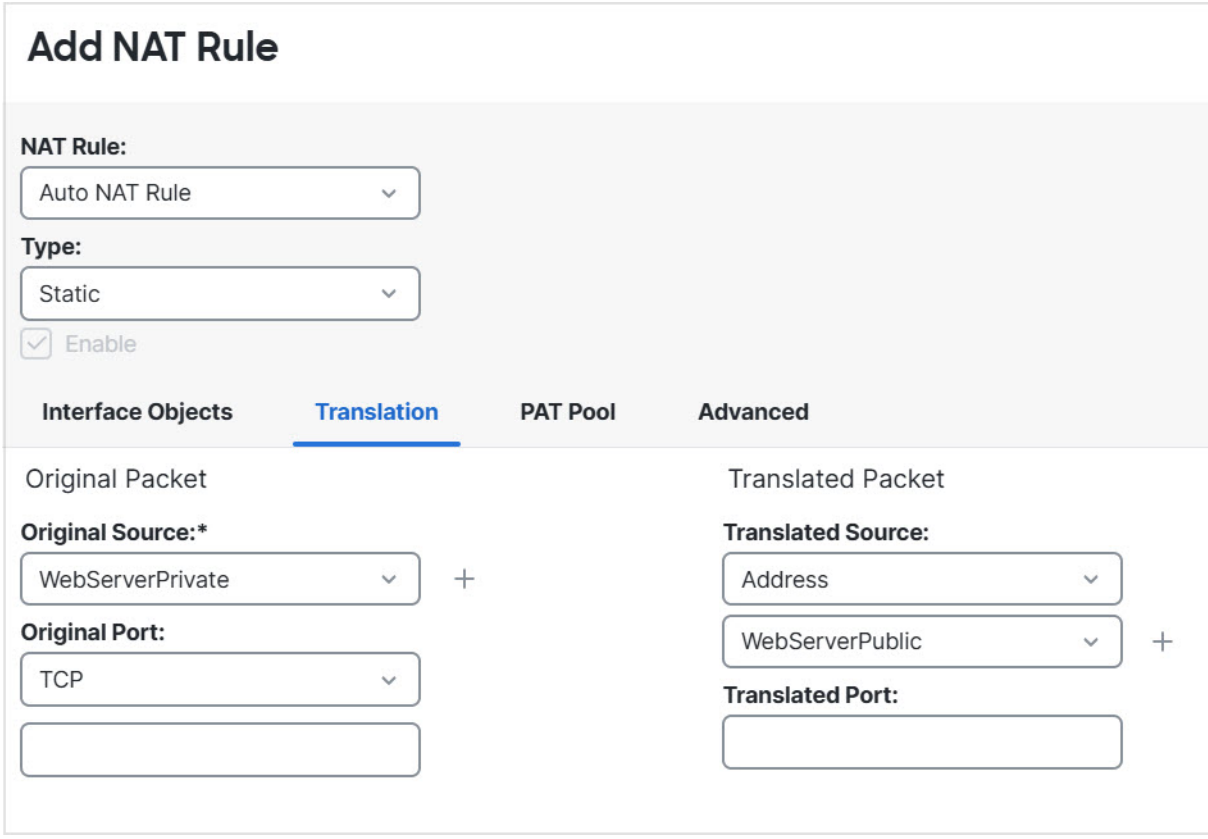

f) Click **Save**.

Step 3 Click **Save** on the NAT rule page.

Dynamic Auto NAT for Inside Hosts and Static NAT for an Outside Web Server

The following example configures dynamic NAT for inside users on a private network when they access the outside. Also, when inside users connect to an outside web server, that web server address is translated to an address that appears to be on the inside network.

Figure 405: Dynamic NAT for Inside, Static NAT for Outside Web Server

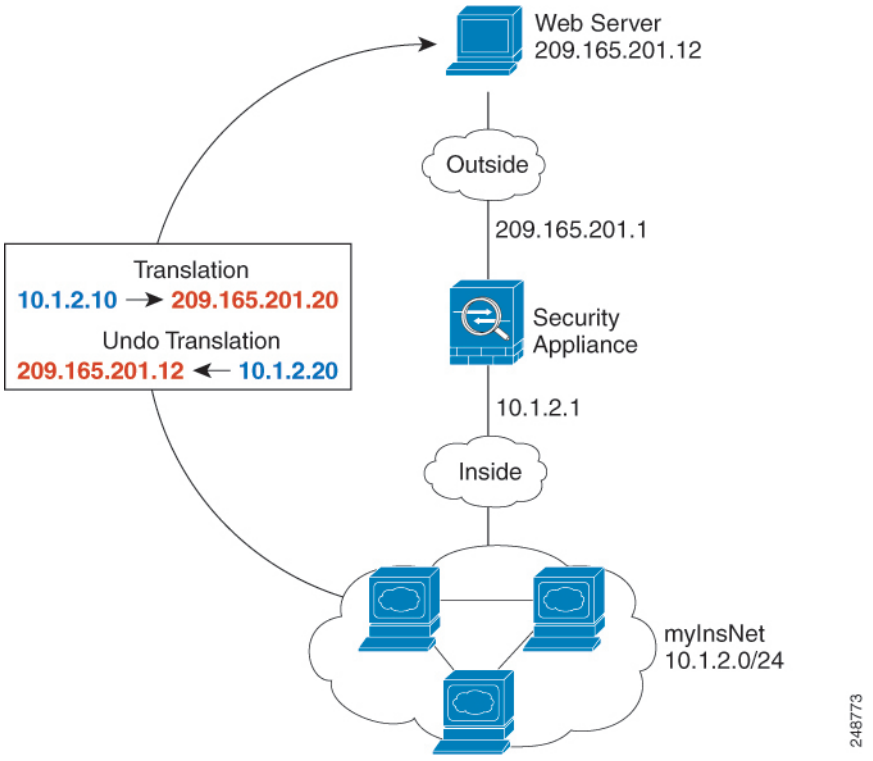

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the web server. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

Step 1 Create a network object for the dynamic NAT pool to which you want to translate the inside addresses.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the dynamic NAT pool.

Name the network object (for example, myNATpool) and enter the network range 209.165.201.20-209.165.201.30.

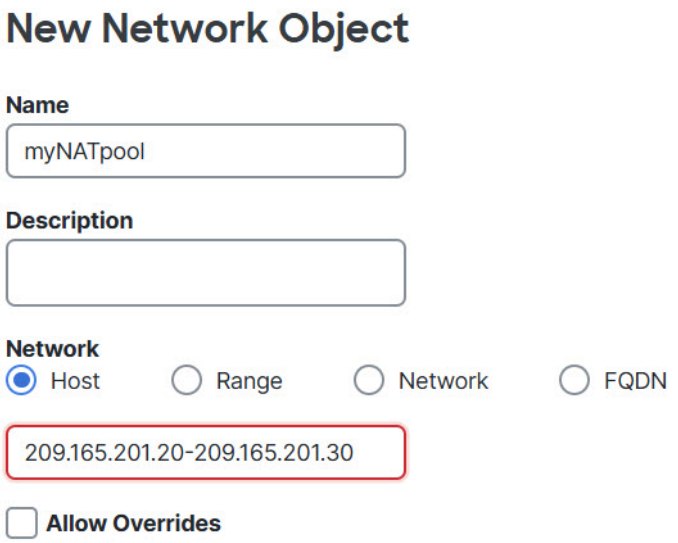

d) Click **Save**.

Step 2 Create a network object for the inside network.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, MyInsNet) and enter the network address 10.1.2.0/24.

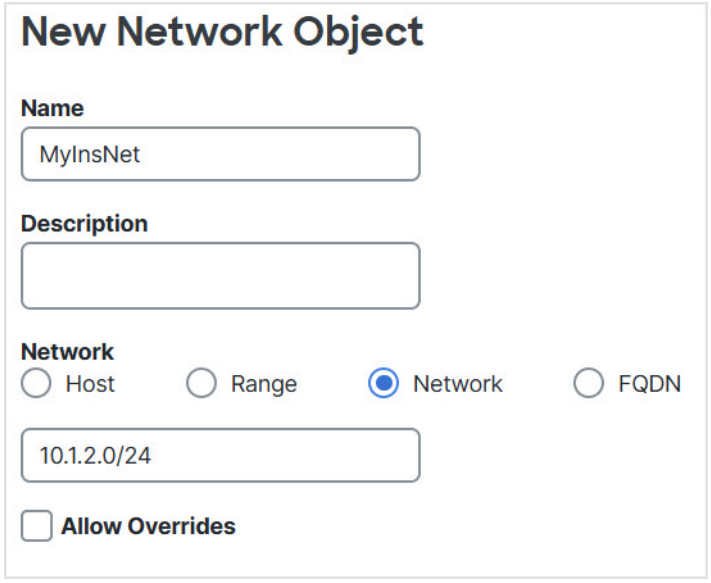

c) Click **Save**.

Step 3 Create a network object for the outside web server.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, MyWebServer) and enter the host address 209.165.201.12.

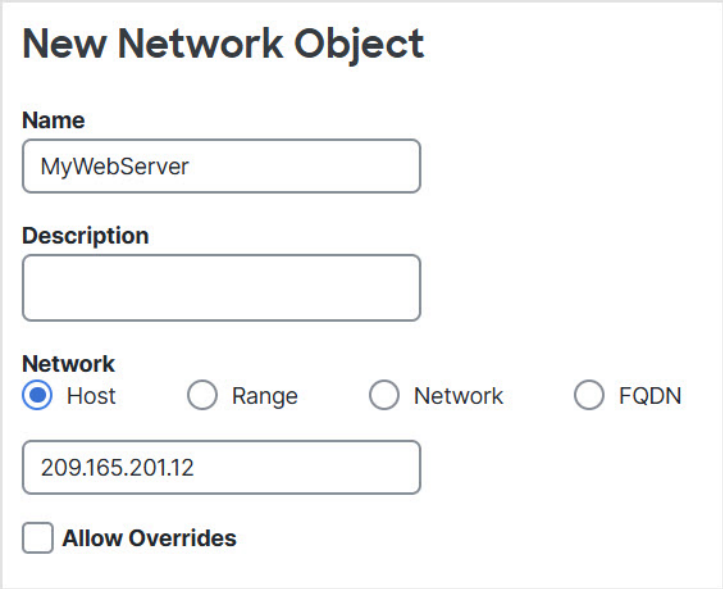

c) Click **Save**.

Step 4 Create a network object for the translated web server address.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, TransWebServer) and enter the host address 10.1.2.20.

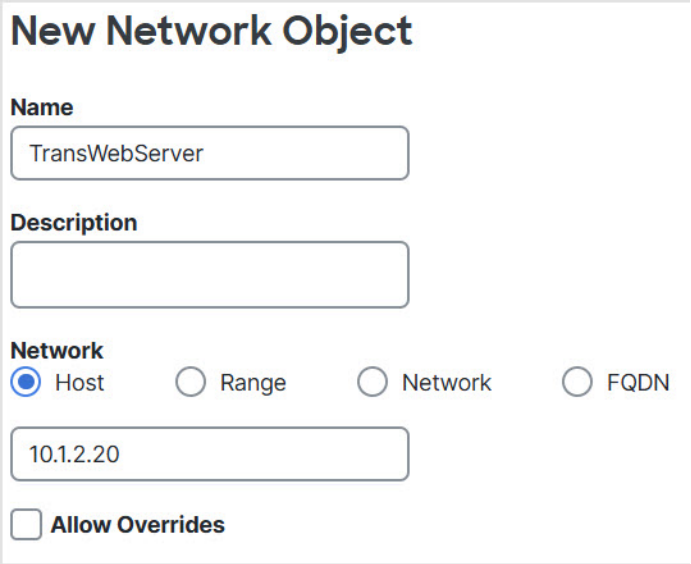

c) Click **Save**.

Step 5 Configure dynamic NAT for the inside network using the dynamic NAT pool object.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.

- **Type** = Dynamic.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = myInsNet network object.
- **Translated Source** > **Address**= myNATpool network group.

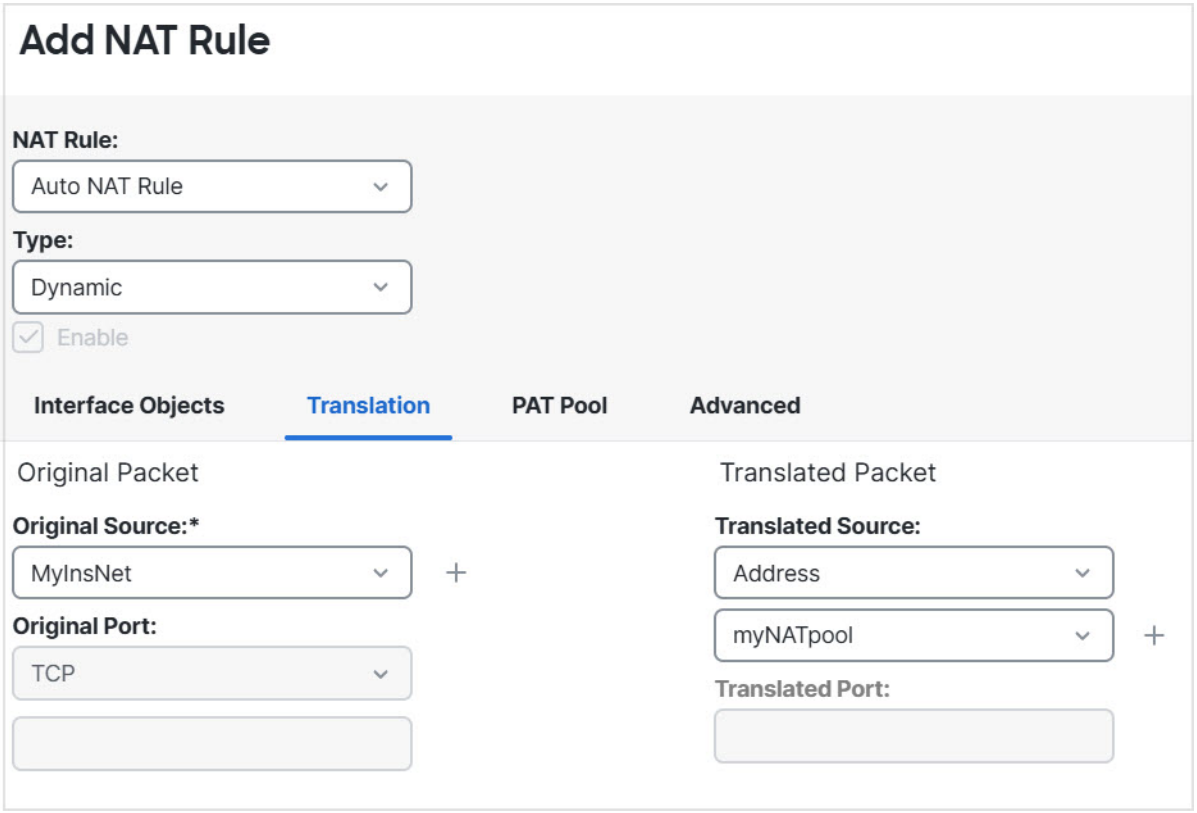

f) Click **Save**.

Step 6 Configure static NAT for the web server.

a) Click **Add Rule**.

b) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.

- **Destination Interface Objects** = inside.

d) On **Translation**, configure the following:

- **Original Source** = myWebServer network object.
- **Translated Source** > **Address**= TransWebServer network object.

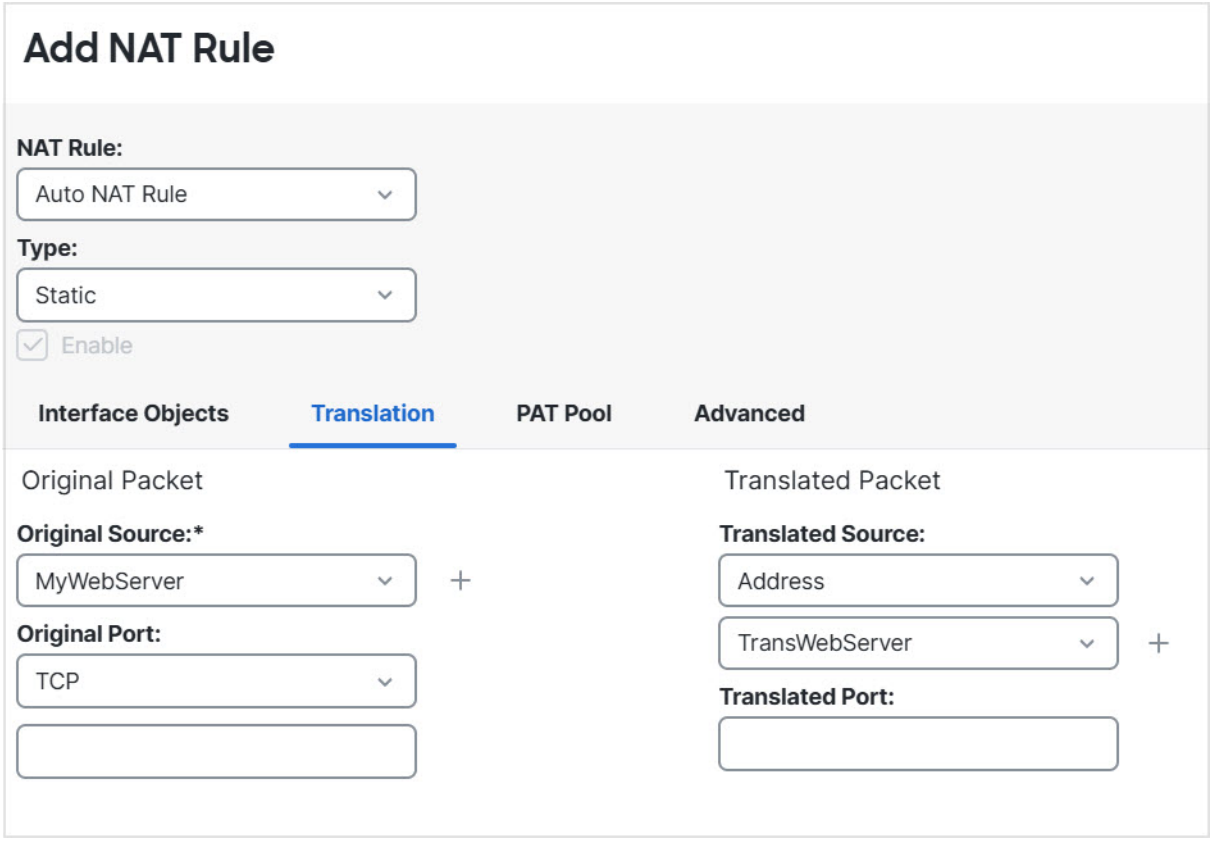

e) Click **Save**.

Step 7 Click **Save** on the NAT rule page.

Inside Load Balancer with Multiple Mapped Addresses (Static Auto NAT, One-to-Many)

The following example shows an inside load balancer that is translated to multiple IP addresses. When an outside host accesses one of the mapped IP addresses, it is untranslated to the single load balancer address. Depending on the URL requested, it redirects traffic to the correct web server.

Figure 406: Static NAT with One-to-Many for an Inside Load Balancer

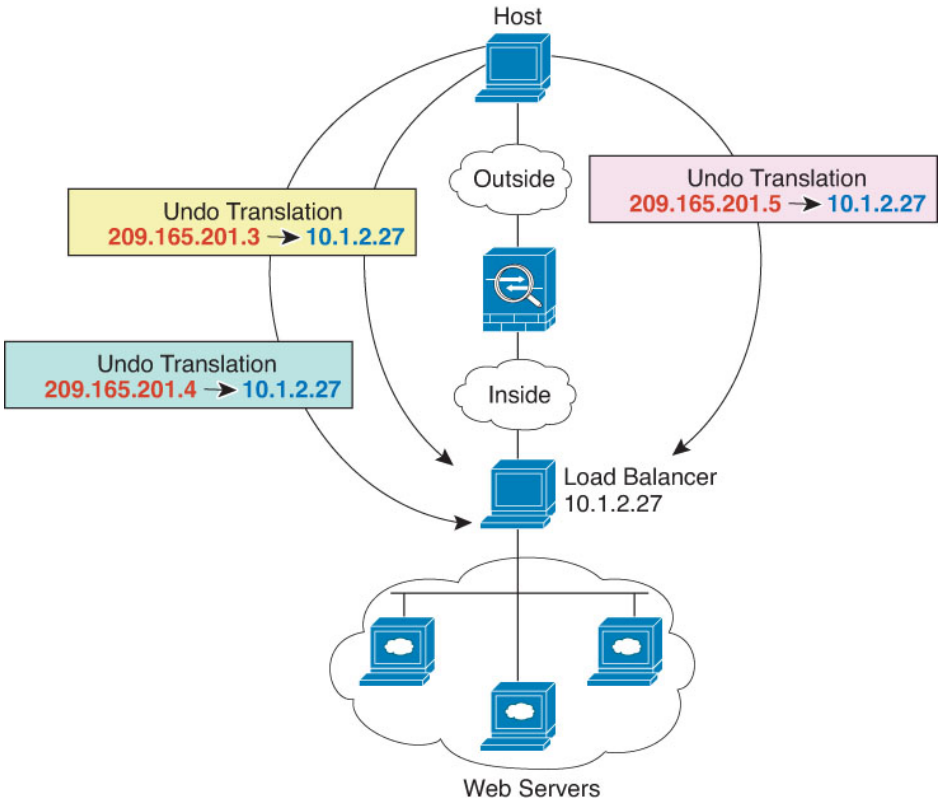

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the web server. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

Step 1 Create a network object for the addresses to which you want to map the load balancer.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the addresses.

Name the network object (for example, myPublicIPs) and enter the network range 209.165.201.3-209.165.201.5.

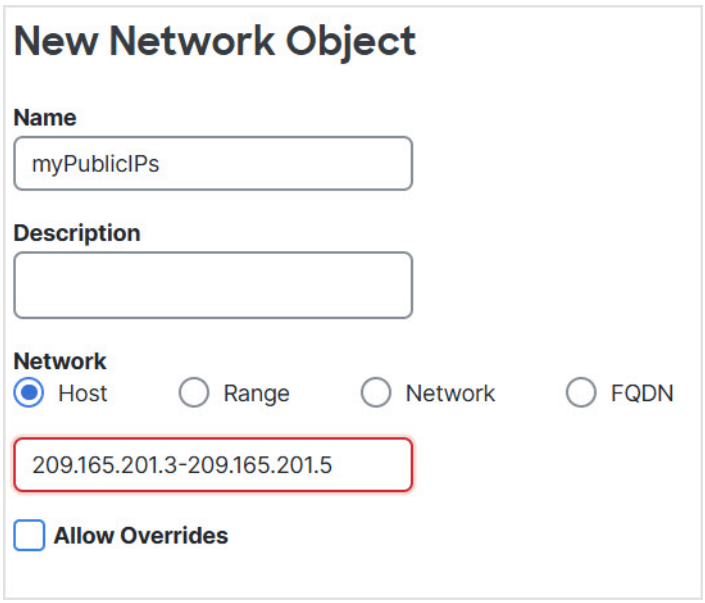

d) Click **Save**.

Step 2 Create a network object for the load balancer.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, myLBHost), enter the host address 10.1.2.27.

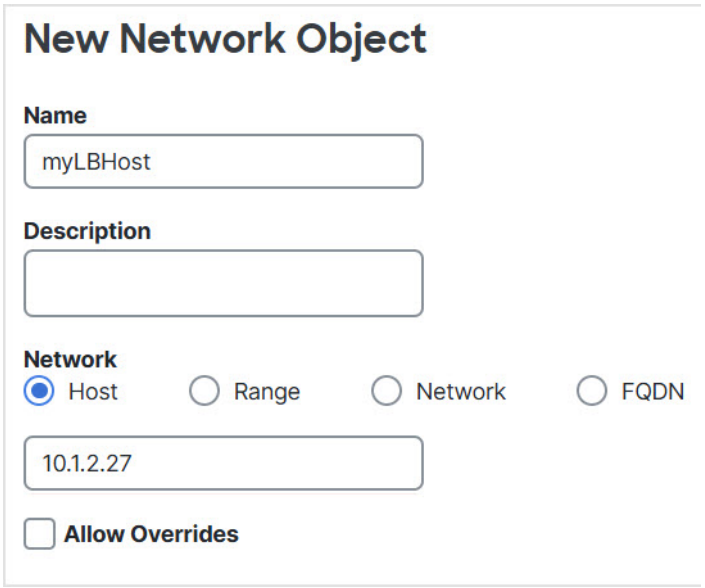

c) Click **Save**.

Step 3 Configure static NAT for the load balancer.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.

- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = myLBHost network object.
- **Translated Source** > **Address**= myPublicIPs network group.

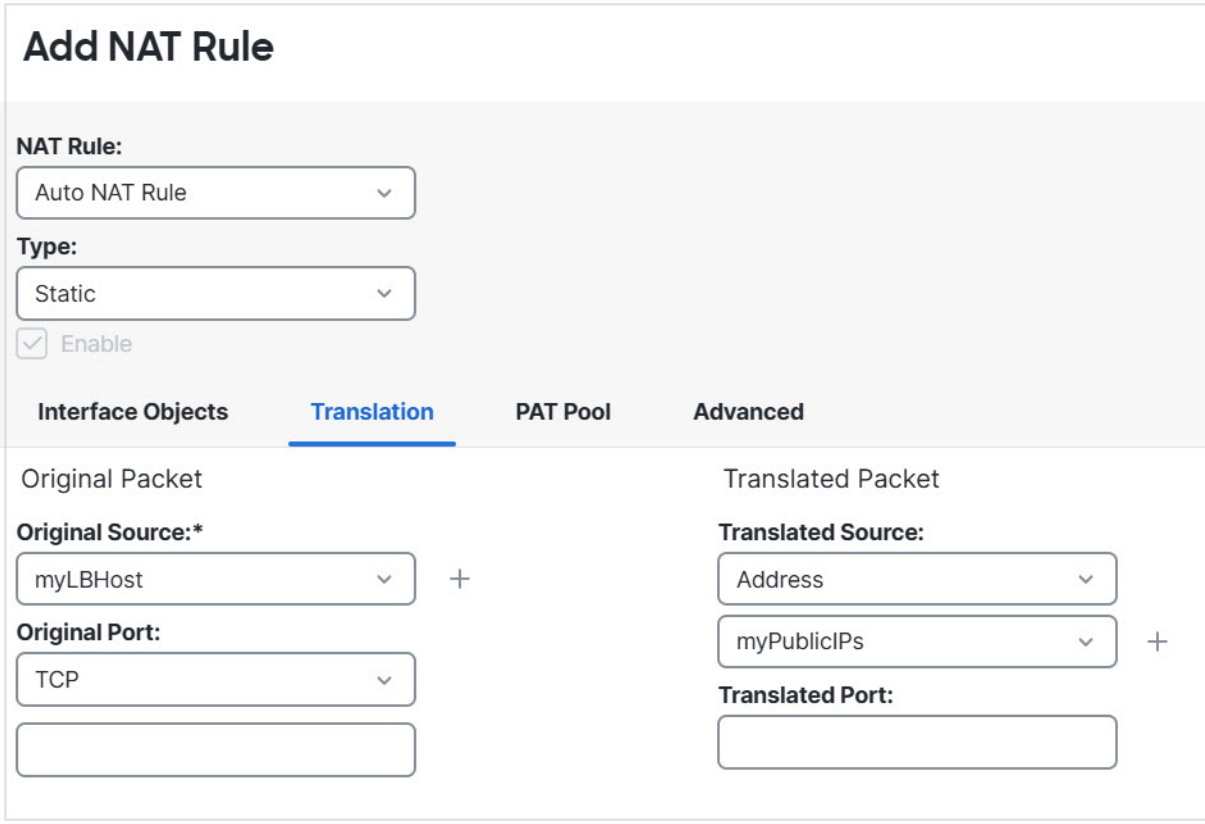

f) Click **Save**.

Step 4 Click **Save** on the NAT rule page.

Single Address for FTP, HTTP, and SMTP (Static Auto NAT-with-Port-Translation)

The following static NAT-with-port-translation example provides a single address for remote users to access FTP, HTTP, and SMTP. These servers are actually different devices on the real network, but for each server,

you can specify static NAT-with-port-translation rules that use the same mapped IP address, but different ports.

Figure 407: Static NAT-with-Port-Translation

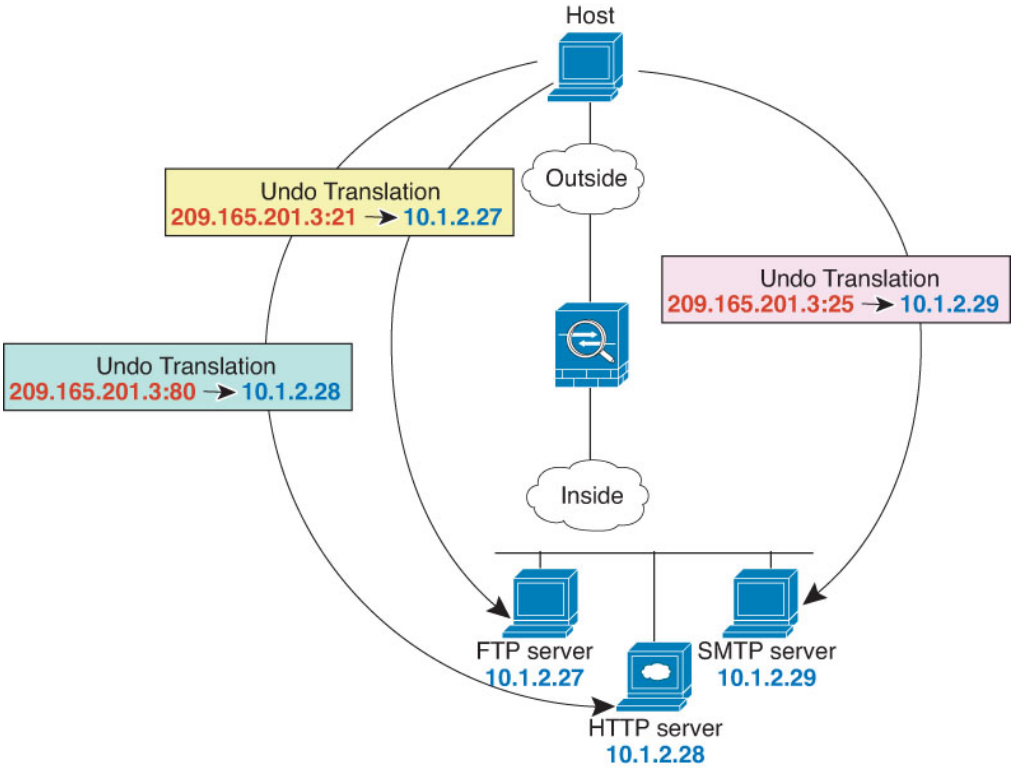

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the servers. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

Step 1 Create a network object for the FTP server.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Name the network object (for example, FTPserver), and enter the real IP address for the FTP server, 10.1.2.27.

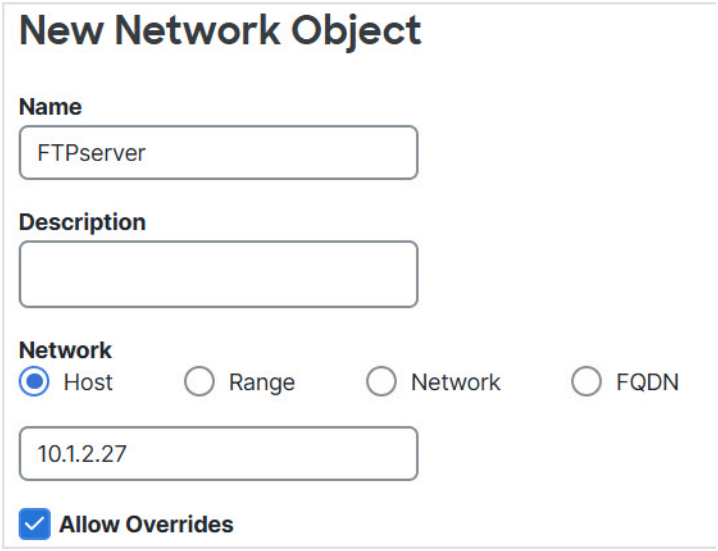

d) Click **Save**.

Step 2 Create a network object for the HTTP server.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, HTTPserver), enter the host address 10.1.2.28.

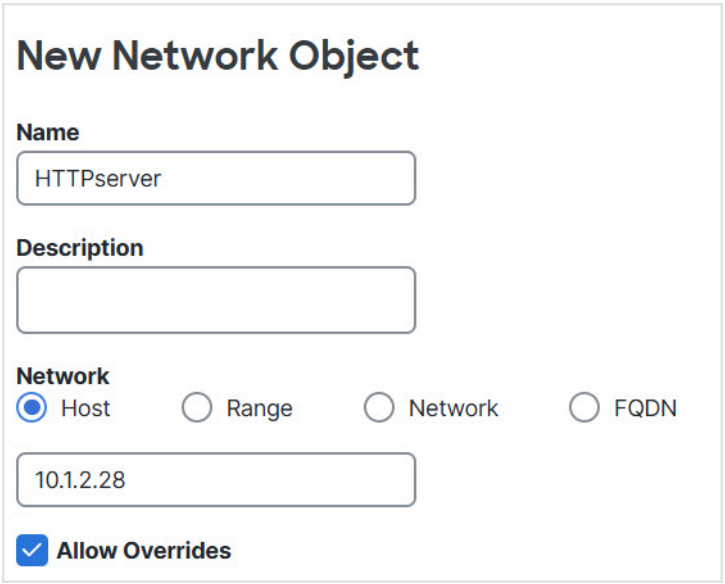

c) Click **Save**.

Step 3 Create a network object for the SMTP server.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, SMTPserver), enter the host address 10.1.2.29.

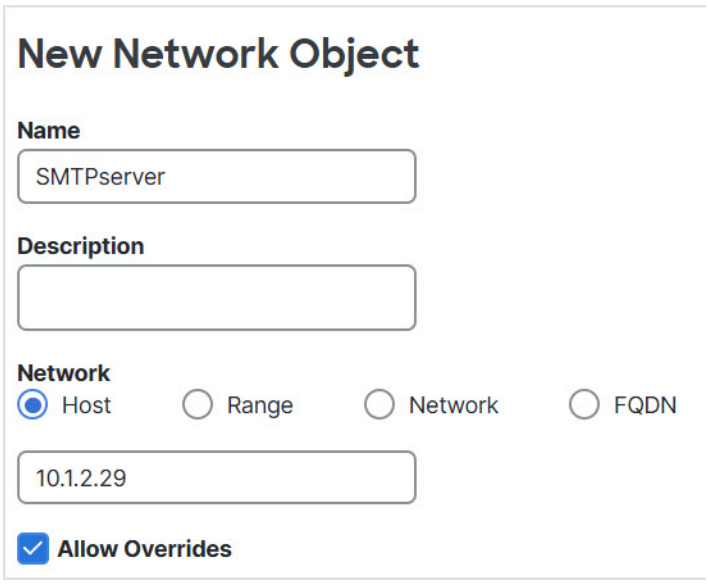

c) Click **Save**.

Step 4 Create a network object for the public IP address used for the three servers.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, ServerPublicIP) and enter the host address 209.165.201.3.

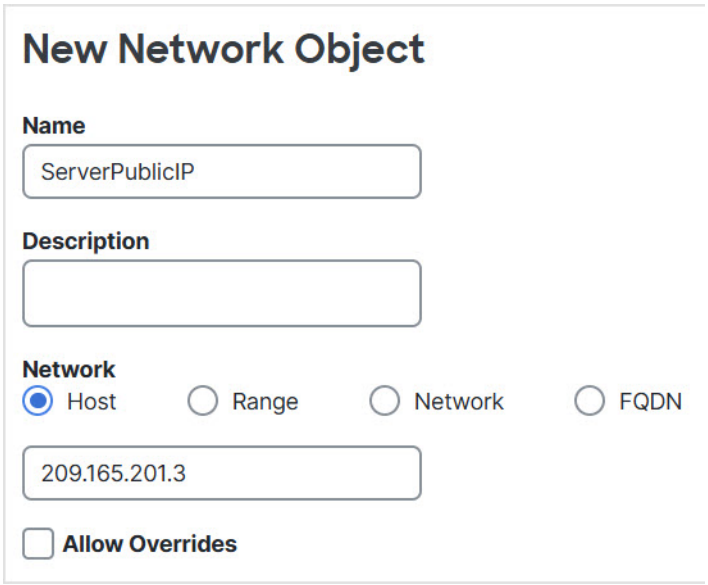

c) Click **Save**.

Step 5 Configure static NAT with port translation for the FTP server, mapping the FTP port to itself.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.

- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = FTPserver network object.
- **Translated Source** > **Address**= ServerPublicIP network object.
- **Original Port** > **TCP** = 21.
- **Translated Port** = 21.

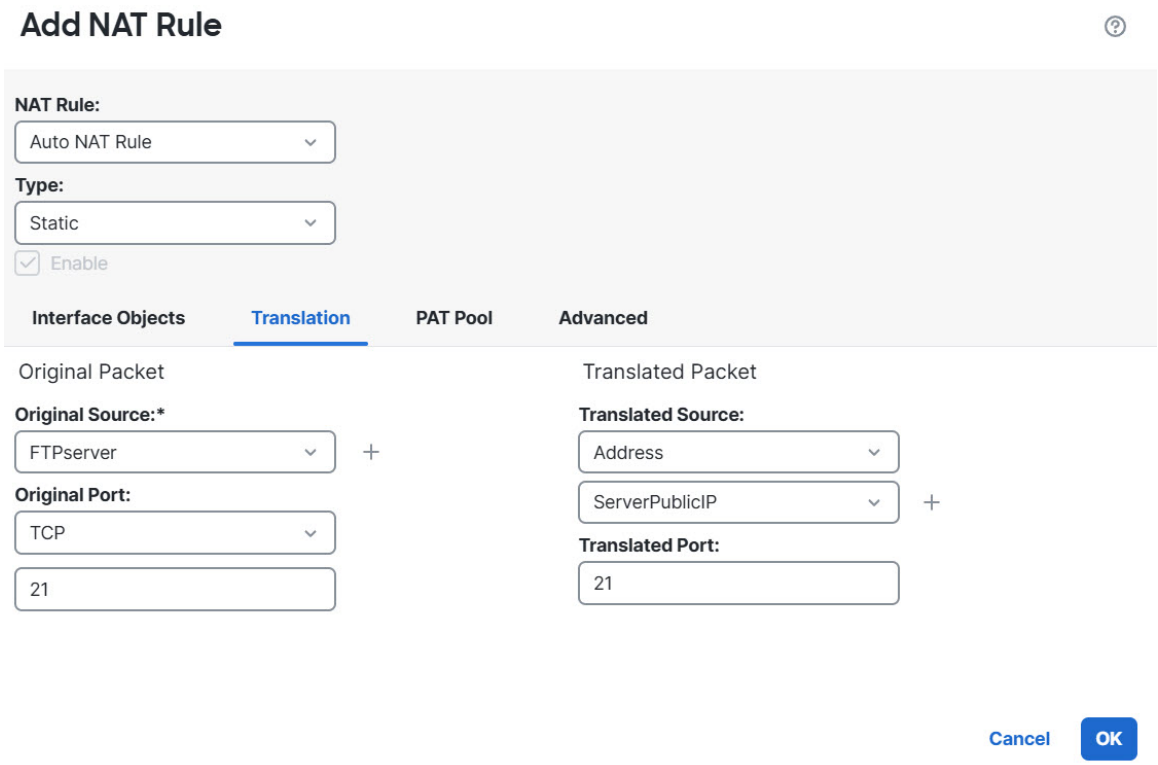

f) Click **Save**.

Step 6 Configure static NAT with port translation for the HTTP server, mapping the HTTP port to itself.

a) Click **Add Rule**.

b) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

d) On **Translation**, configure the following:

- **Original Source** = HTTPserver network object.
- **Translated Source** > **Address**= ServerPublicIP network object.
- **Original Port** > **TCP** = 80.
- **Translated Port** = 80.

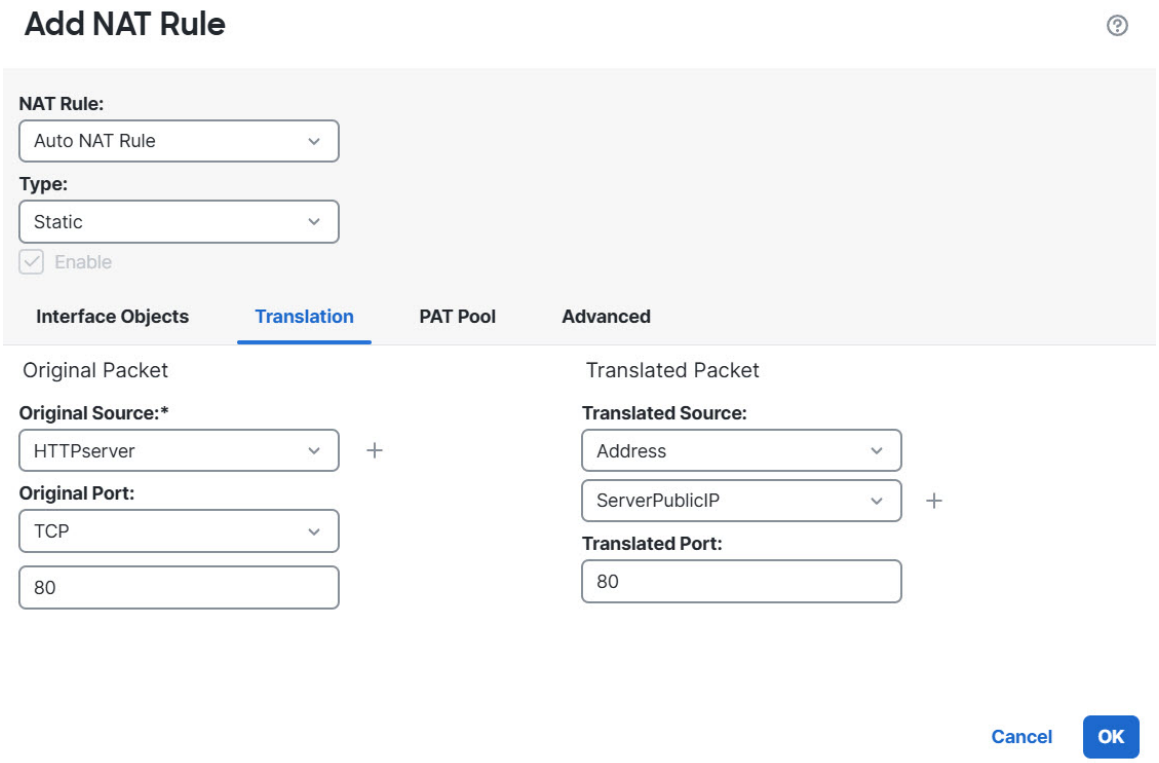

e) Click **Save**.

Step 7 Configure static NAT with port translation for the SMTP server, mapping the SMTP port to itself.

a) Click **Add Rule**.

b) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

d) On **Translation**, configure the following:

- **Original Source** = SMTPserver network object.
- **Translated Source** > **Address**= ServerPublicIP network object.
- **Original Port** > **TCP** = 25.
- **Translated Port** = 25.

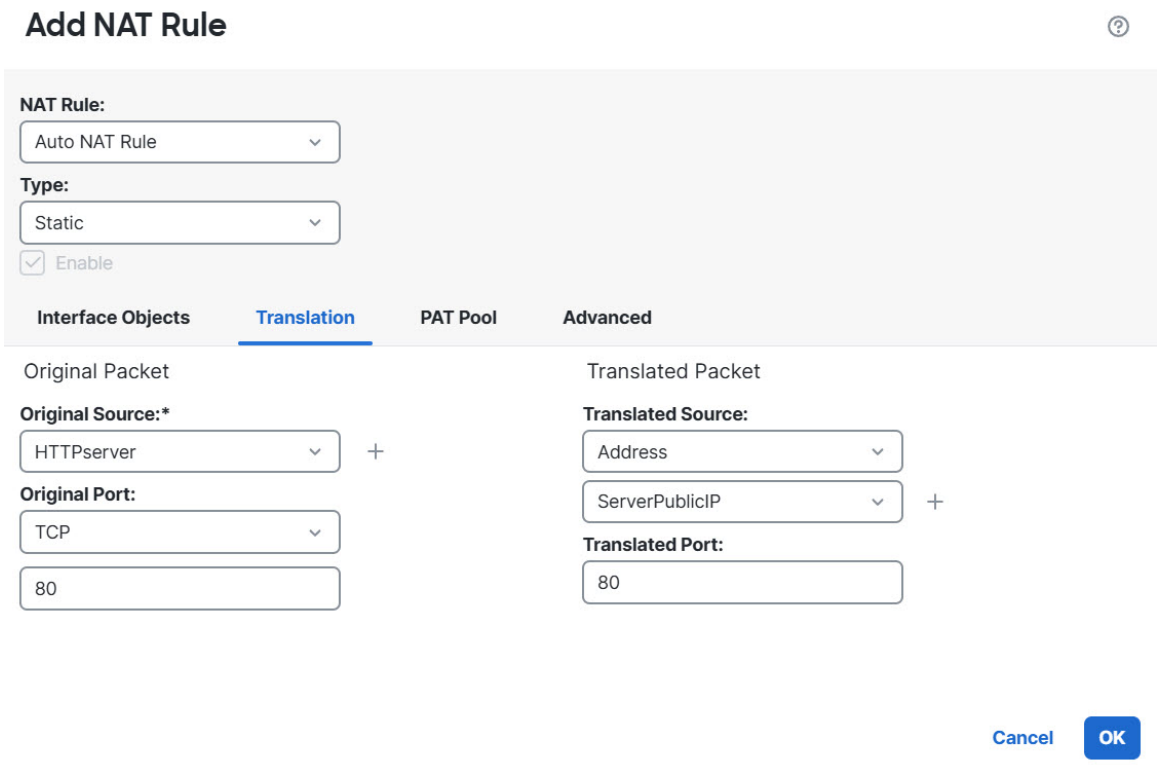

e) Click **Save**.

Step 8 Click **Save** on the NAT rule page.

Different Translation Depending on the Destination (Dynamic Manual PAT)

The following figure shows a host on the 10.1.2.0/24 network accessing two different servers. When the host accesses the server at 209.165.201.11, the real address is translated to 209.165.202.129:*port*. When the host accesses the server at 209.165.200.225, the real address is translated to 209.165.202.130:*port*.

Figure 408: Manual NAT with Different Destination Addresses

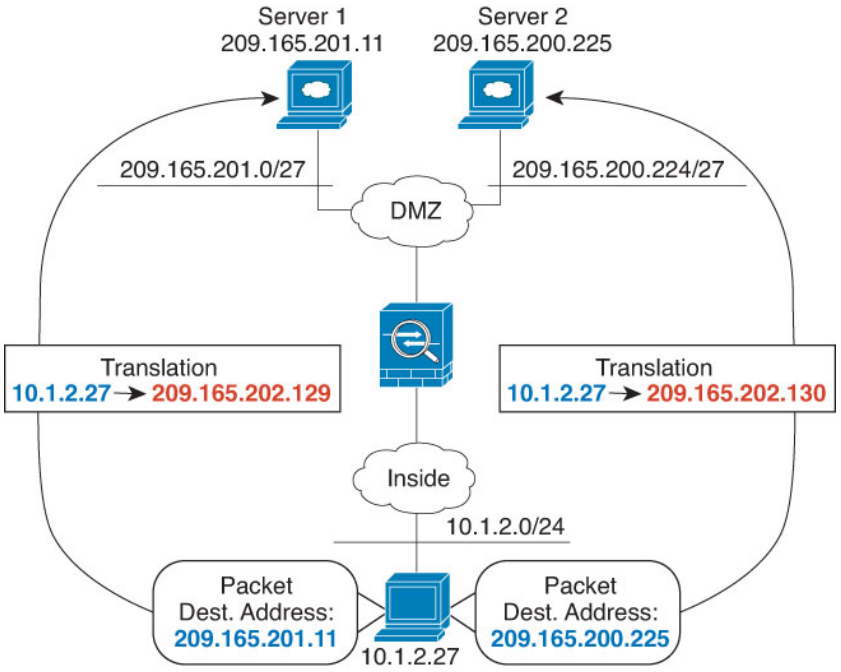

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the servers. In this example, we will assume the interface objects are security zones named **inside** and **dmz**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

Step 1 Create a network object for the inside network.

a) Choose **Objects** > **Object Management**.
b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.
c) Name the network object (for example, myInsideNetwork), and enter the real network address, 10.1.2.0/24.

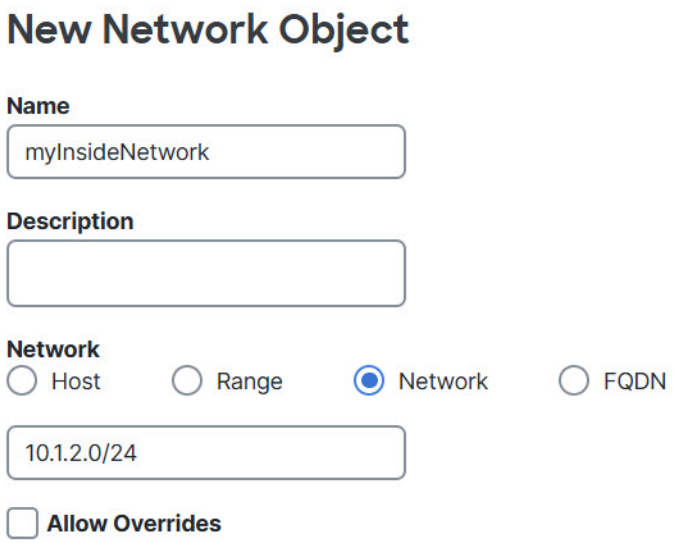

d) Click **Save**.

Step 2 Create a network object for the DMZ network 1.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, DMZnetwork1) and enter the network address 209.165.201.0/27 (subnet mask of 255.255.255.224).

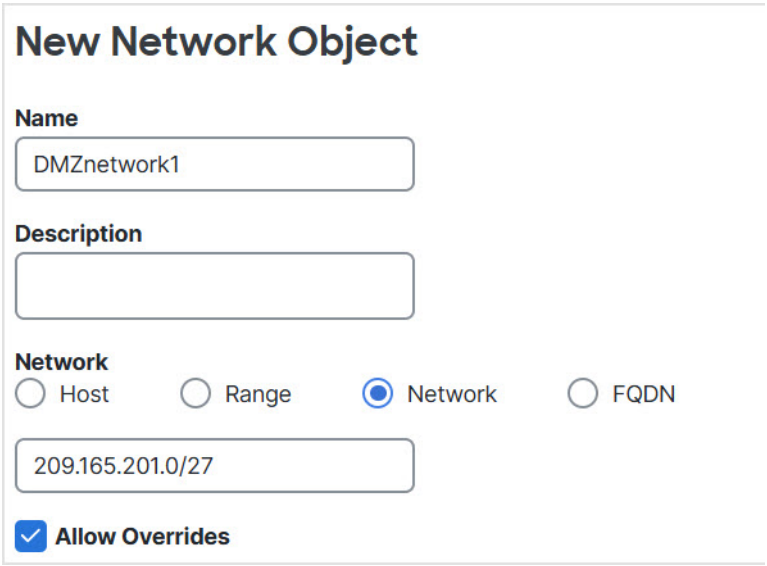

c) Click **Save**.

Step 3 Create a network object for the PAT address for DMZ network 1.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, PATaddress1) and enter the host address 209.165.202.129.

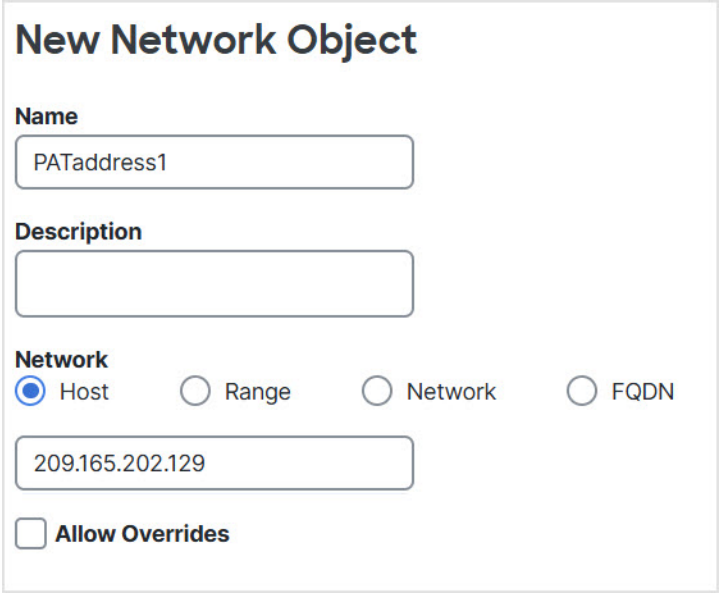

c) Click **Save**.

Step 4 Create a network object for the DMZ network 2.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, DMZnetwork2) and enter the network address 209.165.200.224/27 (subnet mask of 255.255.255.224).

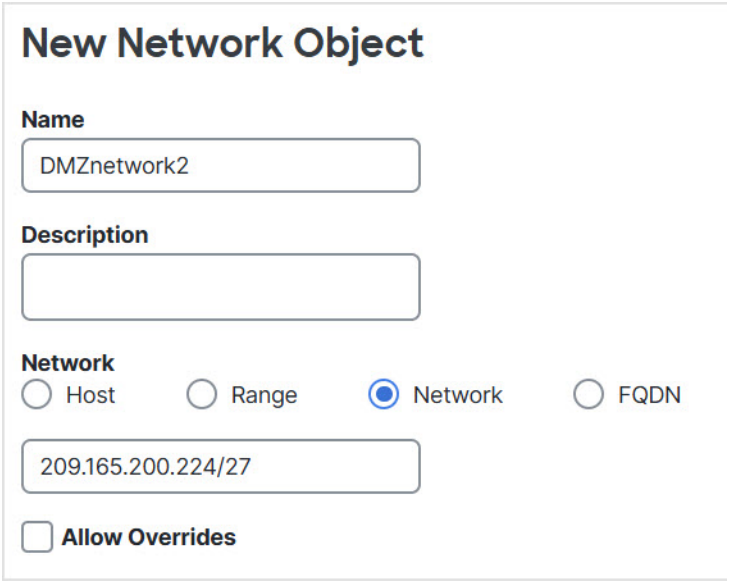

c) Click **Save**.

Step 5 Create a network object for the PAT address for DMZ network 2.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, PATaddress2) and enter the host address 209.165.202.130.

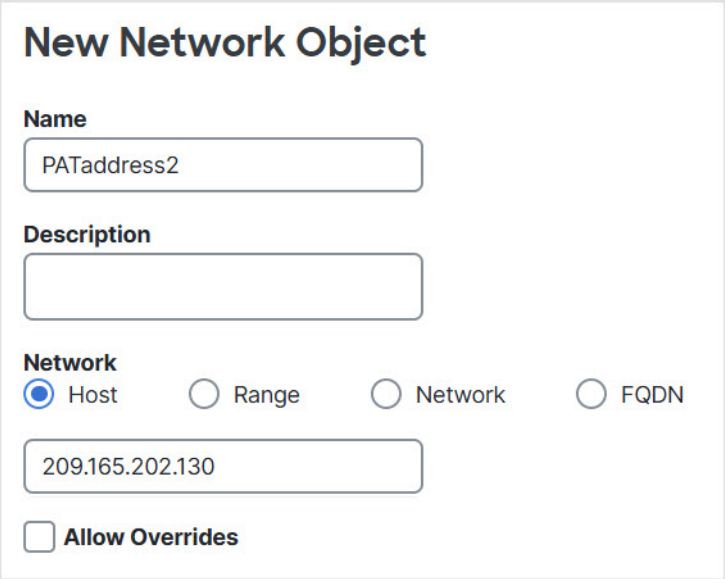

c) Click **Save**.

Step 6 Configure dynamic manual PAT for DMZ network 1.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = dmz.

e) On **Translation**, configure the following:

- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address**= PATaddress1 network object.
- **Original Destination** > **Address** = DMZnetwork1 network object.
- **Translated Destination** = DMZnetwork1 network object.

Note
Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the original and translated destination addresses. Leave all of the port fields blank.

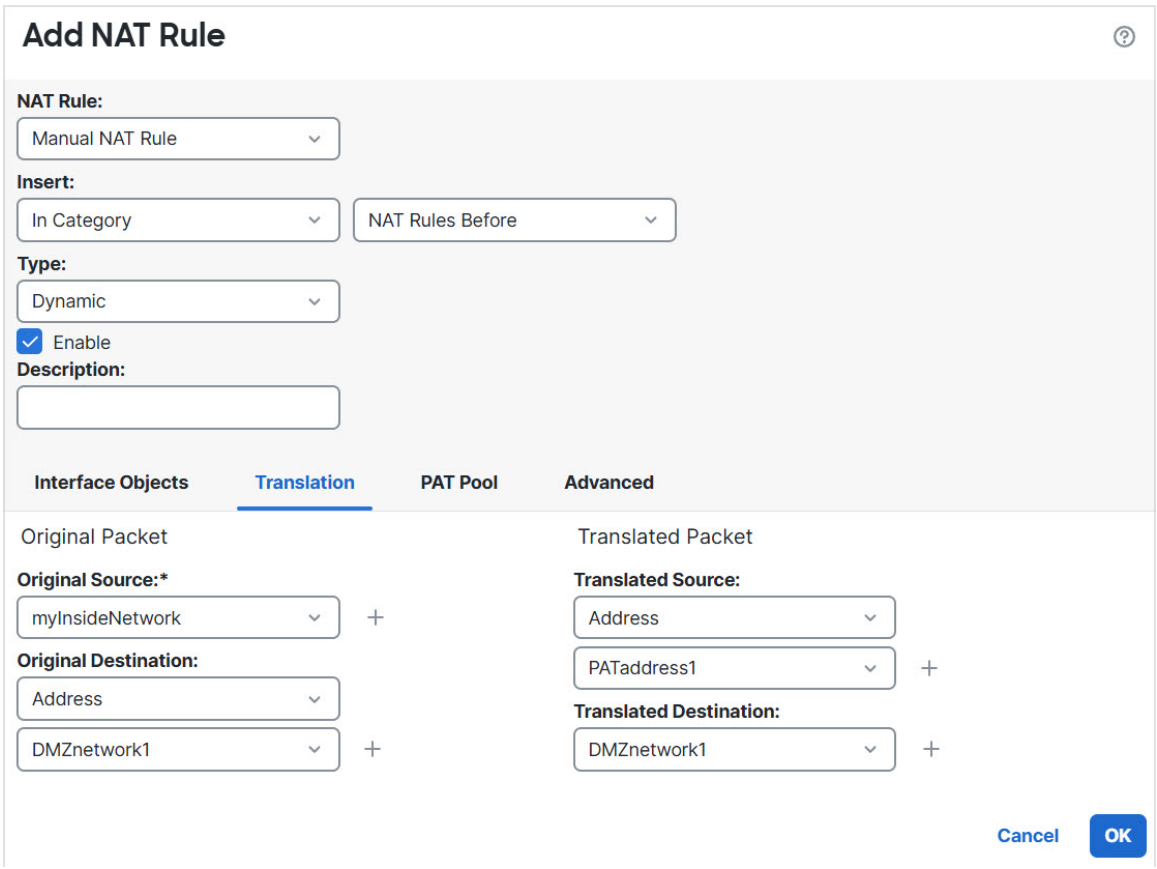

f) Click **Save**.

Step 7 Configure dynamic manual PAT for DMZ network 2.

a) Click **Add Rule**.

b) Configure the following properties:

- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = dmz.

d) On **Translation**, configure the following:

- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address**= PATaddress2 network object.
- **Original Destination** > **Address** = DMZnetwork2 network object.
- **Translated Destination** = DMZnetwork2 network object.

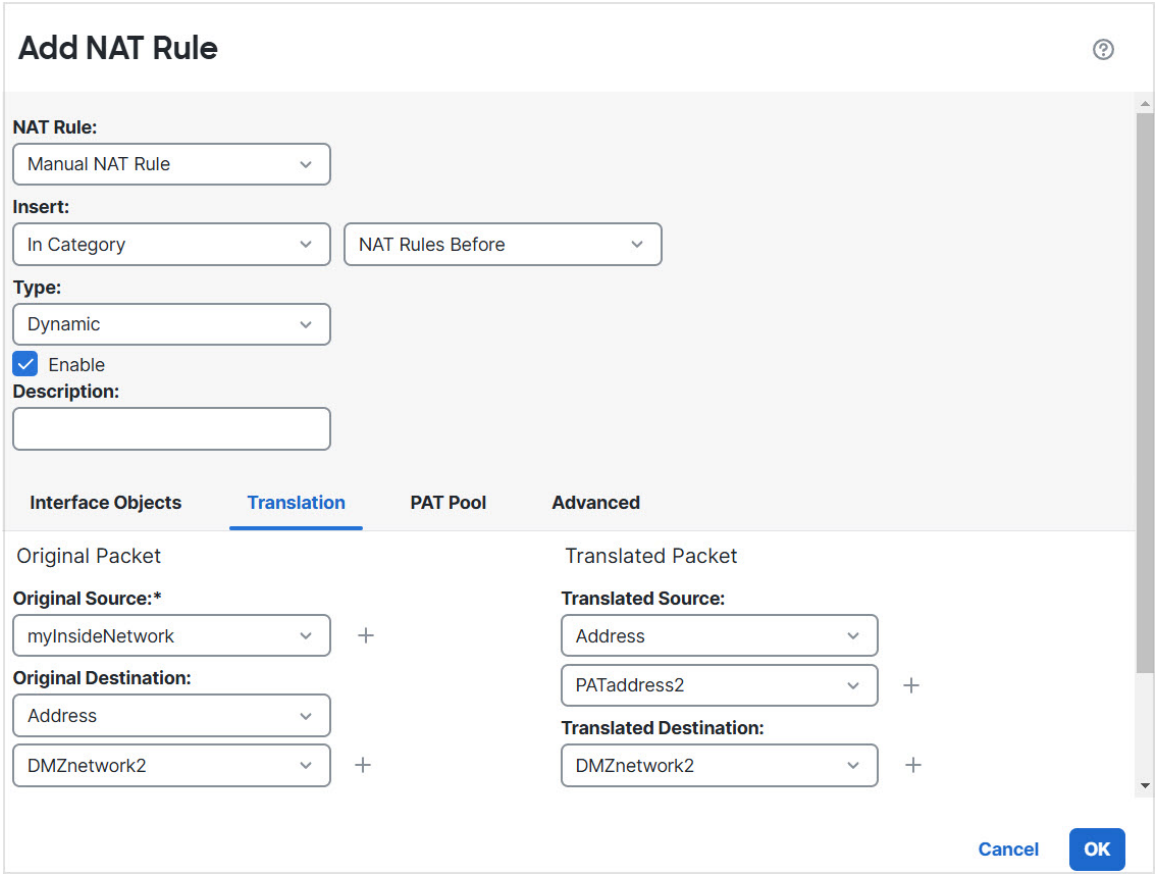

e) Click **Save**.

Step 8 Click **Save** on the NAT rule page.

Different Translation Depending on the Destination Address and Port (Dynamic Manual PAT)

The following figure shows the use of source and destination ports. The host on the 10.1.2.0/24 network accesses a single host for both web services and Telnet services. When the host accesses the server for Telnet services, the real address is translated to 209.165.202.129:*port*. When the host accesses the same server for web services, the real address is translated to 209.165.202.130:*port*.

Figure 409: Manual NAT with Different Destination Ports

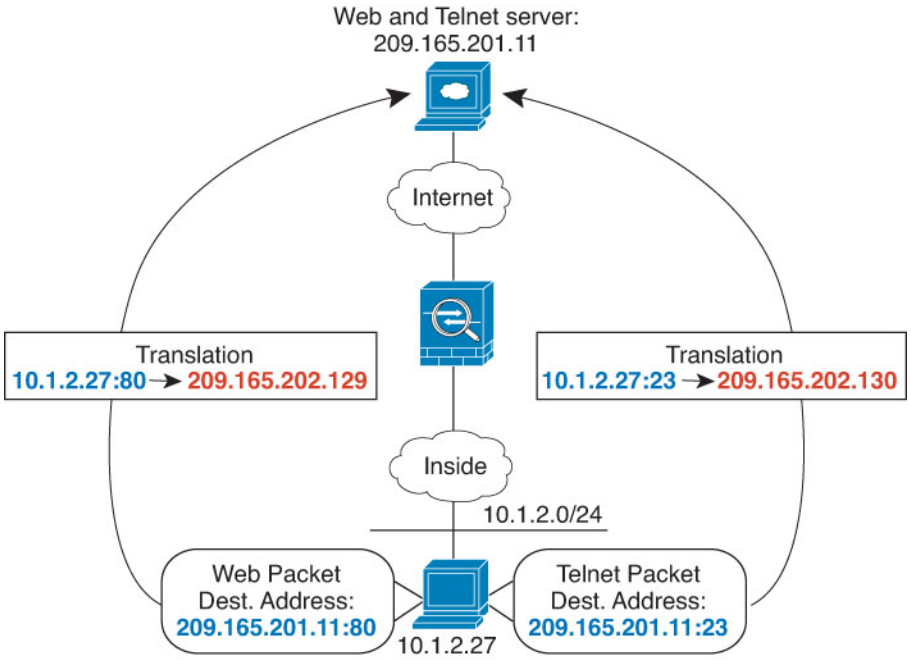

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device that protects the servers. In this example, we will assume the interface objects are security zones named **inside** and **dmz**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

Step 1 Create a network object for the inside network.

a) Choose **Objects** > **Object Management**.
b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.
c) Name the network object (for example, myInsideNetwork) and enter the real network address, 10.1.2.0/24.

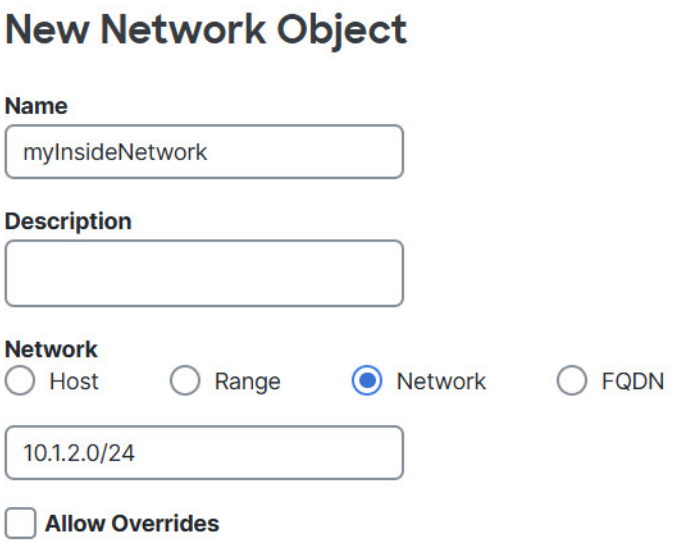

d) Click **Save**.

Step 2 Create a network object for the Telnet/Web server.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, TelnetWebServer) and enter the host address 209.165.201.11.

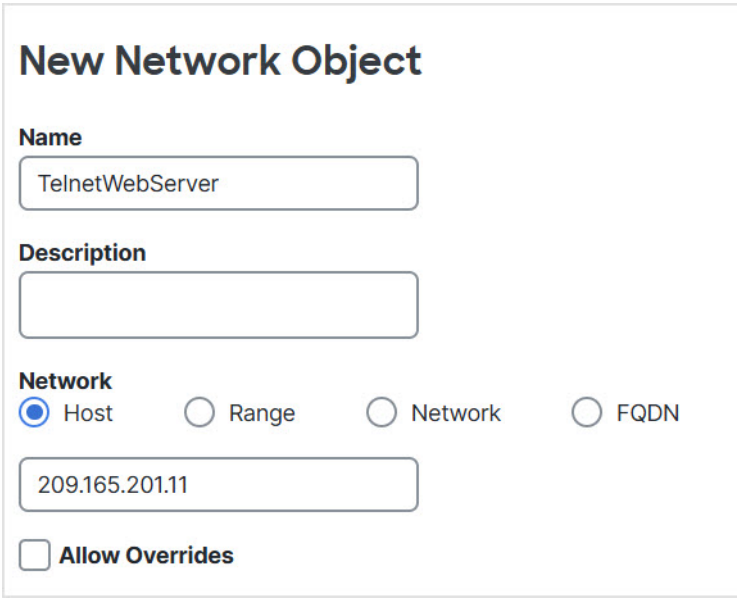

c) Click **Save**.

Step 3 Create a network object for the PAT address when using Telnet.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, PATaddress1) and enter the host address 209.165.202.129.

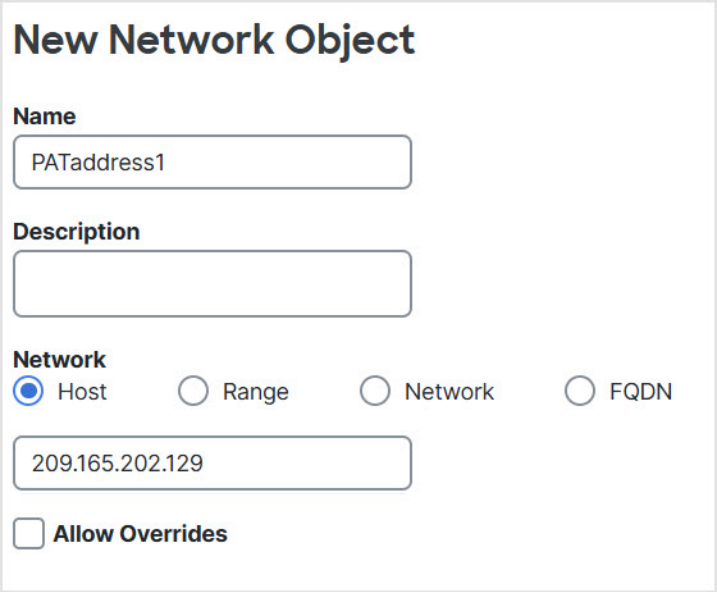

c) Click **Save**.

Step 4 Create a network object for the PAT address when using HTTP.

a) Click **Add Network** > **Add Object**.

b) Name the network object (for example, PATaddress2) and enter the host address 209.165.202.130.

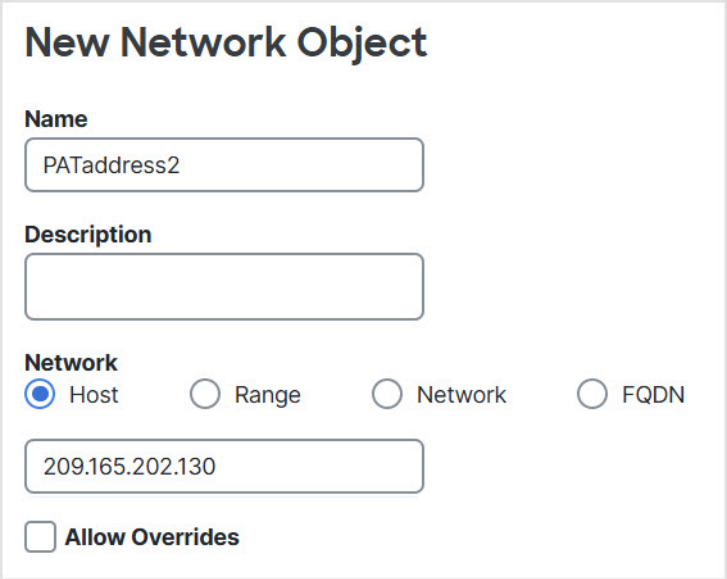

c) Click **Save**.

Step 5 Configure dynamic manual PAT for Telnet access.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Manual NAT Rule.

- **Type** = Dynamic.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = dmz.

e) On **Translation**, configure the following:

- **Original Source** = myInsideNetwork network object.
- **Translated Source** > **Address**= PATaddress1 network object.
- **Original Destination** > **Address** = TelnetWebServer network object.
- **Translated Destination** = TelnetWebServer network object.
- **Original Destination Port** = TELNET port object (system-defined).
- **Translated Destination Port** = TELNET port object (system-defined).

Note
Because you do not want to translate the destination address or port, you need to configure identity NAT for them by specifying the same address for the original and translated destination addresses, and the same port for the original and translated port.

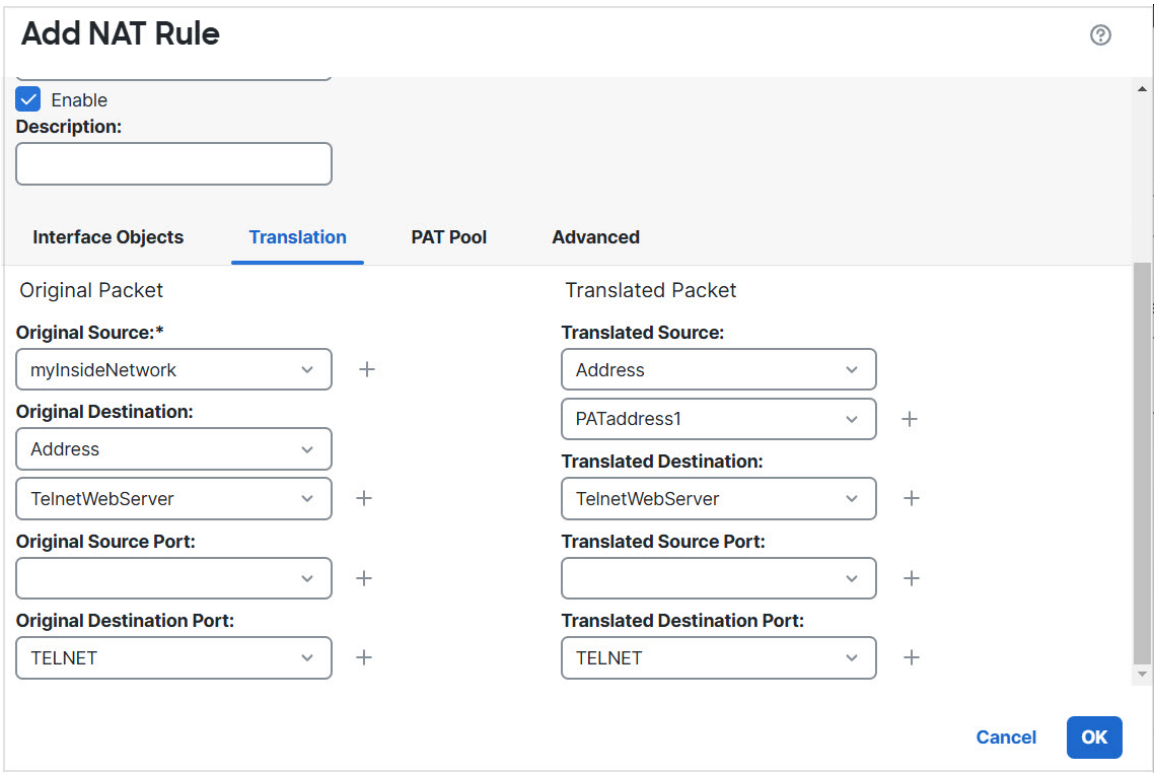

f) Click **Save**.

Step 6 Configure dynamic manual PAT for web access.

a) Click **Add Rule**.
b) Configure the following properties:
   - **NAT Rule** = Manual NAT Rule.
   - **Type** = Dynamic.
c) On **Interface Objects**, configure the following:
   - **Source Interface Objects** = inside.
   - **Destination Interface Objects** = dmz.
d) On **Translation**, configure the following:
   - **Original Source** = myInsideNetwork network object.
   - **Translated Source** > **Address**= PATaddress2 network object.
   - **Original Destination** > **Address** = TelnetWebServer network object.
   - **Translated Destination** = TelnetWebServer network object.
   - **Original Destination Port** = HTTP port object (system-defined).
   - **Translated Destination Port** = HTTP port object (system-defined).

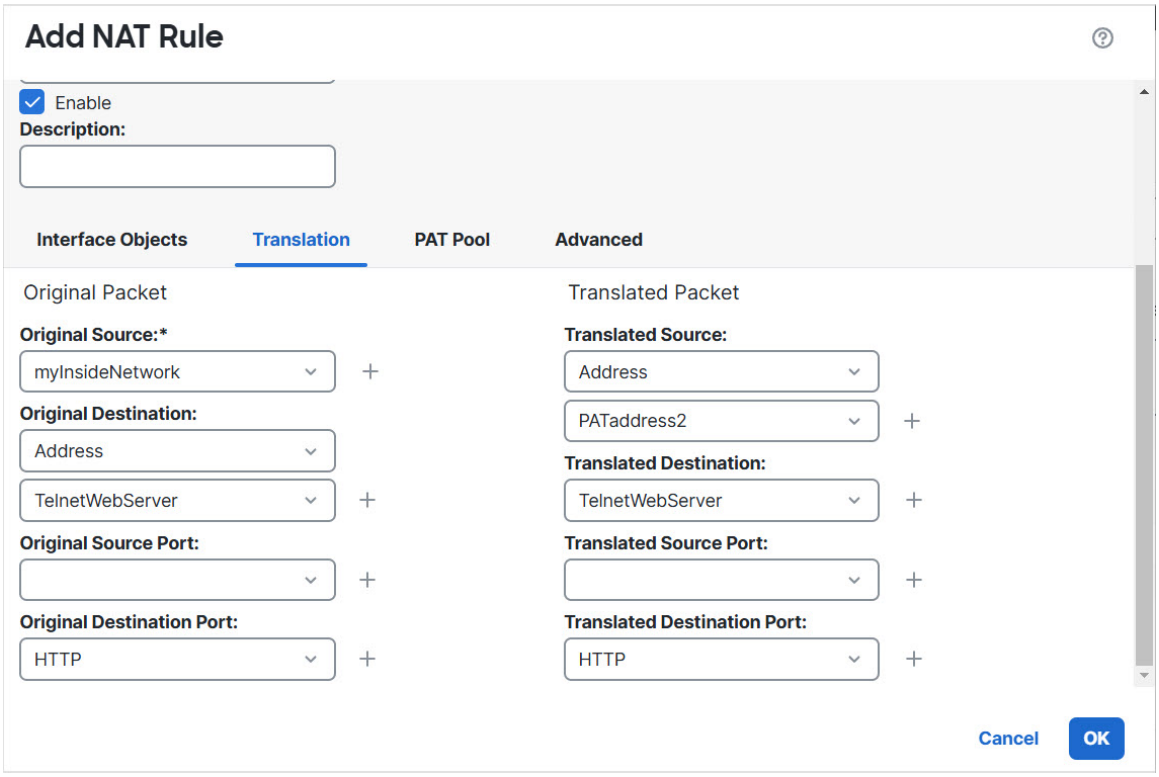

e) Click **Save**.

Step 7 Click **Save** on the NAT rule page.

NAT and Site-to-Site VPN

When you create a policy-based site-to-site VPN using the management center VPN wizard (**Device > Site To Site**), you can select the NAT Exempt option to create the rules automatically. You can view the NAT exemptions for a device in the NAT policy page (**Device > NAT > NAT Exemptions**).If you do not want to configure NAT Exempt in the VPN wizard, you can use the following procedure for NAT exemption.

The following figure shows a site-to-site tunnel connecting the Boulder and San Jose offices. For traffic that you want to go to the Internet (for example from 10.1.1.6 in Boulder to www.example.com), you need a public IP address provided by NAT to access the Internet. The below example uses interface PAT rules. However, for traffic that you want to go over the VPN tunnel (for example from 10.1.1.6 in Boulder to 10.2.2.78 in San Jose), you do not want to perform NAT; you need to exempt that traffic by creating an identity NAT rule. Identity NAT simply translates an address to the same address.

Figure 410: Interface PAT and Identity NAT for Site-to-Site VPN

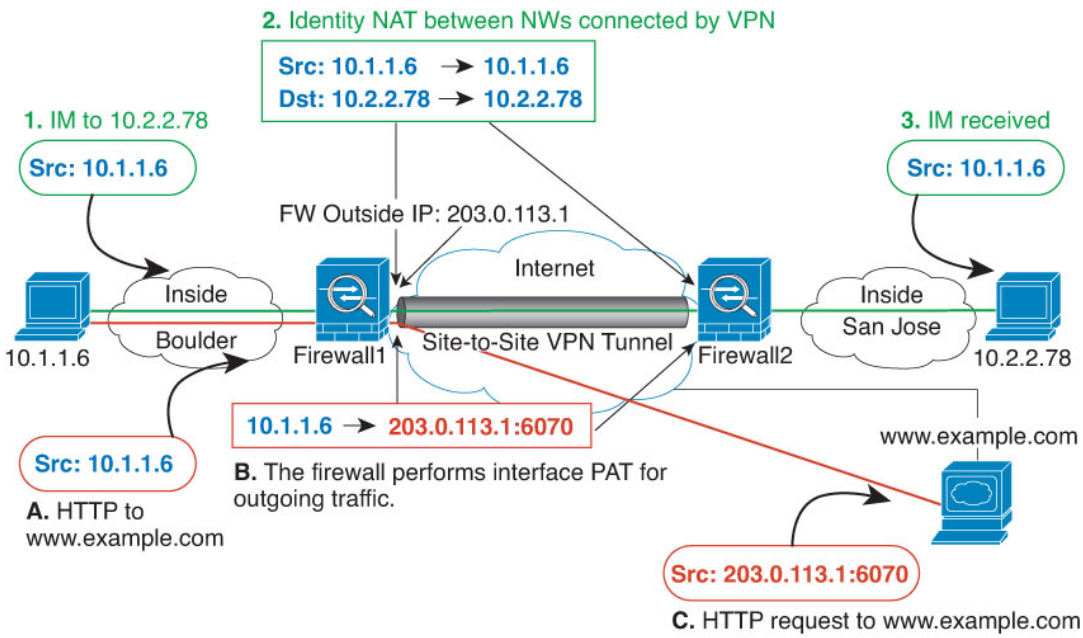

The following example explains the configuration for Firewall1 (Boulder).

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the devices in the VPN. In this example, we will assume the interface objects are security zones named **inside-boulder** and **outside-boulder** for the Firewall1 (Boulder) interfaces. To configure interface objects, select **Objects** > **Object Management**, then select **Interfaces**.

Procedure

Step 1 Create the objects to define the various networks.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Identify the Boulder inside network.

Name the network object (for example, boulder-network) and enter the network address, 10.1.1.0/24.

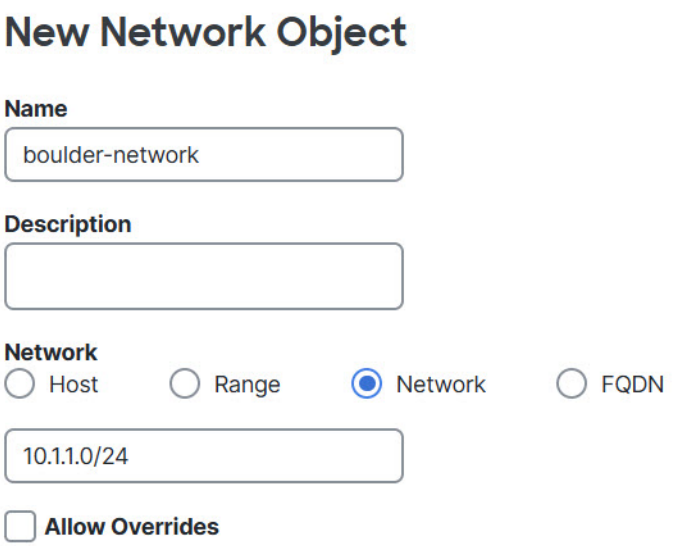

d) Click **Save**.

e) Click **Add Network** > **Add Object** and define the inside San Jose network.

Name the network object (for example, sanjose-network) and enter the network address 10.2.2.0/24.

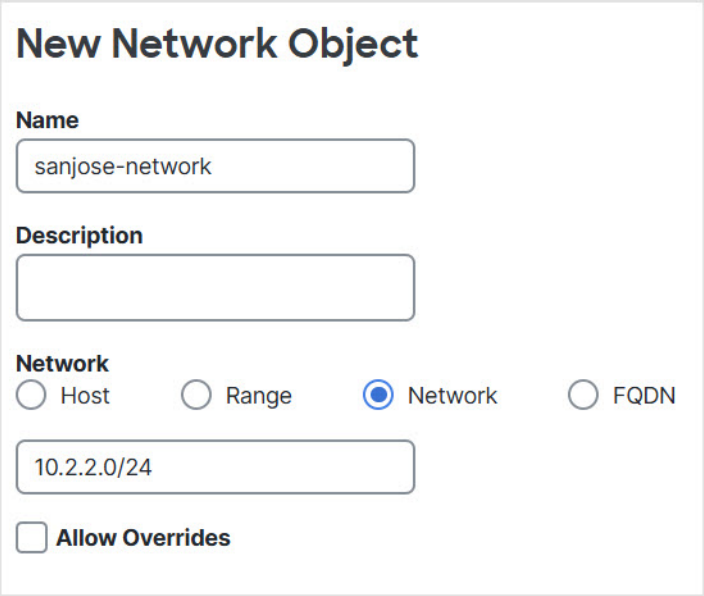

f) Click **Save**.

Step 2 Configure manual identity NAT for the Boulder network when going over the VPN to San Jose on Firewall1 (Boulder).

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.
b) Click **Add Rule**.
c) Configure the following properties:

- **NAT Rule** = Manual NAT Rule.
- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside-boulder.
- **Destination Interface Objects** = outside-boulder.

e) On **Translation**, configure the following:

- **Original Source** = boulder-network object.
- **Translated Source** > **Address** = boulder-network object.
- **Original Destination** > **Address** = sanjose-network object.
- **Translated Destination** = sanjose-network object.

Note
Because you do not want to translate the destination address, you need to configure identity NAT for it by specifying the same address for the original and translated destination addresses. Leave all of the port fields blank. This rule configures identity NAT for both source and destination.

f) On **Advanced**, select **Do not proxy ARP on Destination interface**.

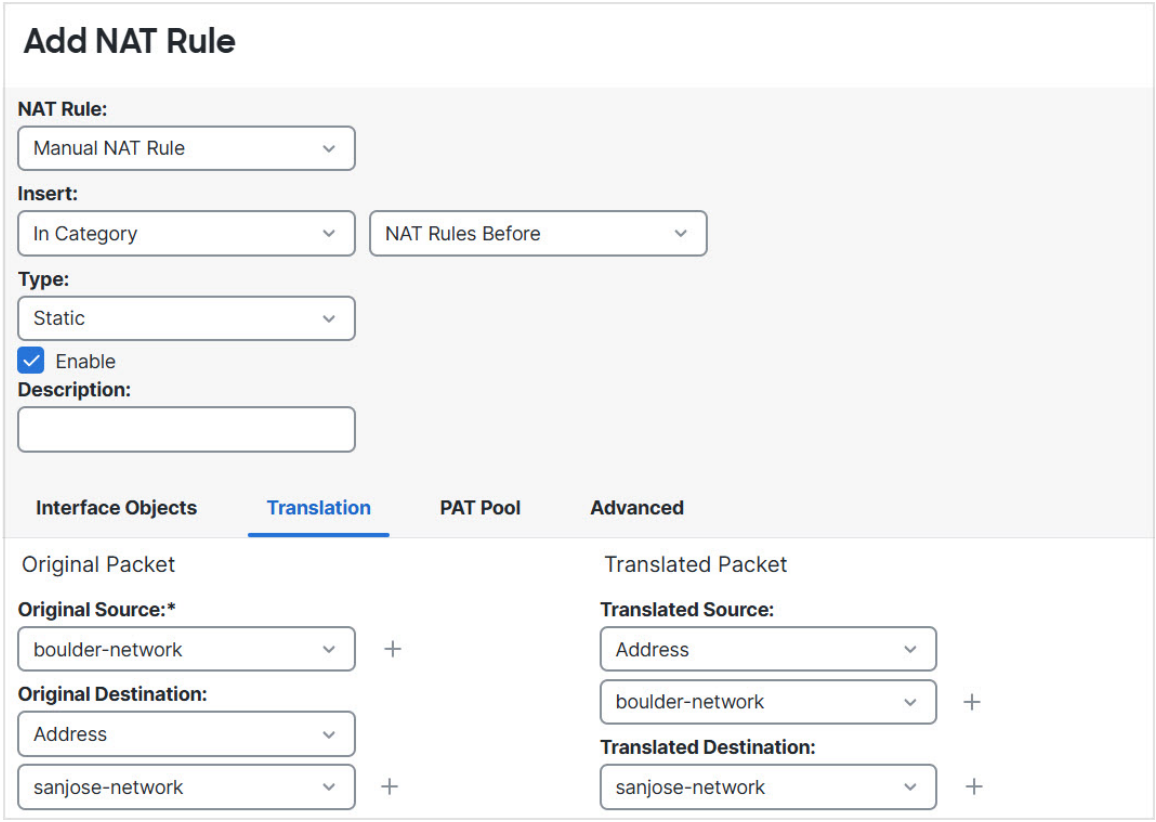

g) Click **Save**.

Step 3 Configure manual dynamic interface PAT when going to the Internet for the inside Boulder network on Firewall1 (Boulder).

a) Click **Add Rule**.

b) Configure the following properties:

- **NAT Rule** = Manual NAT Rule.
- **Type** = Dynamic.
- **Insert Rule** = any position after the first rule. Because this rule will apply to any destination address, the rule that uses sanjose-network as the destination must come before this rule, or the sanjose-network rule will never be matched. The default is to place new manual NAT rules at the end of the "NAT Rules Before Auto NAT" section.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside-boulder.
- **Destination Interface Objects** = outside-boulder.

d) On **Translation**, configure the following:

- **Original Source** = boulder-network object.
- **Translated Source** = **Destination Interface IP**. This option configures interface PAT using the interface contained in the destination interface object.

- **Original Destination** > **Address** = any (leave blank).
- **Translated Destination** = any (leave blank).

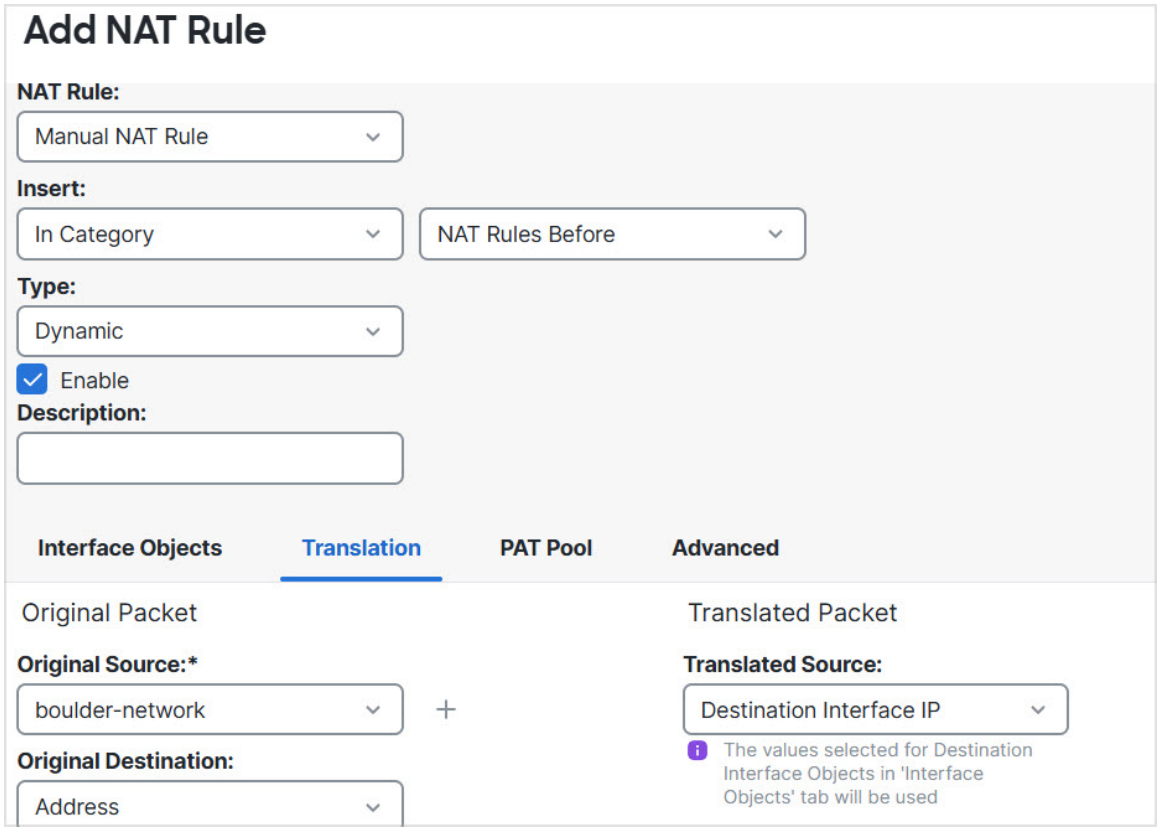

e) Click **Save**.

Step 4 If you are also managing Firewall2 (San Jose), you can configure similar rules for that device.

- The manual identity NAT rule would be for sanjose-network when the destination is boulder-network. Create new interface objects for the Firewall2 inside and outside networks.
- The manual dynamic interface PAT rule would be for sanjose-network when the destination is "any."

Rewriting DNS Queries and Responses Using NAT

You might need to configure the Firewall Threat Defense device to modify DNS replies by replacing the address in the reply with an address that matches the NAT configuration. You can configure DNS modification when you configure each translation rule. DNS modification is also known as DNS doctoring.

This feature rewrites the address in DNS queries and replies that match a NAT rule (for example, the A record for IPv4, the AAAA record for IPv6, or the PTR record for reverse DNS queries). For DNS replies traversing from a mapped interface to any other interface, the record is rewritten from the mapped value to the real value. Inversely, for DNS replies traversing from any interface to a mapped interface, the record is rewritten from the real value to the mapped value. This feature works with NAT44,NAT 66, NAT46, and NAT64.

Following are the main circumstances when you would need to configure DNS rewrite on a NAT rule.

- The rule is NAT64 or NAT46, and the DNS server is on the outside network. You need DNS rewrite to convert between DNS A records (for IPv4) and AAAA records (for IPv6).
- The DNS server is on the outside, clients are on the inside, and some of the fully-qualified domain names that the clients use resolve to other inside hosts.
- The DNS server is on the inside and responds with private IP addresses, clients are on the outside, and the clients access fully-qualified domain names that point to servers that are hosted on the inside.

DNS Rewrite Limitations

Following are some limitations with DNS rewrite:

- DNS rewrite is not applicable for PAT because multiple PAT rules are applicable for each A or AAAA record, and the PAT rule to use is ambiguous.
- If you configure a manual NAT rule, you cannot configure DNS modification if you specify the destination address as well as the source address. These kinds of rules can potentially have a different translation for a single address when going to A vs. B. Therefore, the can not accurately match the IP address inside the DNS reply to the correct twice NAT rule; the DNS reply does not contain information about which source/destination address combination was in the packet that prompted the DNS request.
- You must enable DNS application inspection with DNS NAT rewrite enabled for NAT rules to rewrite DNS queries and responses. By default, DNS inspection with DNS NAT rewrite enabled is globally applied, so you probably do not need to change the inspection configuration.
- DNS rewrite is actually done on the xlate entry, not the NAT rule. Thus, if there is no xlate for a dynamic rule, rewrite cannot be done correctly. The same problem does not occur for static NAT.
- DNS rewrite does not rewrite DNS Dynamic Update messages (opcode 5).

The following topics provide examples of DNS rewrite in NAT rules.

DNS64 Reply Modification

The following figure shows an FTP server and DNS server on the outside IPv4 network. The system has a static translation for the outside server. In this case, when an inside IPv6 user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.200.225.

Because you want inside users to use the mapped address for ftp.cisco.com (2001:DB8::D1A5:C8E1, where D1A5:C8E1 is the IPv6 equivalent of 209.165.200.225) you need to configure DNS reply modification for the static translation. This example also includes a static NAT translation for the DNS server, and a PAT rule for the inside IPv6 hosts.

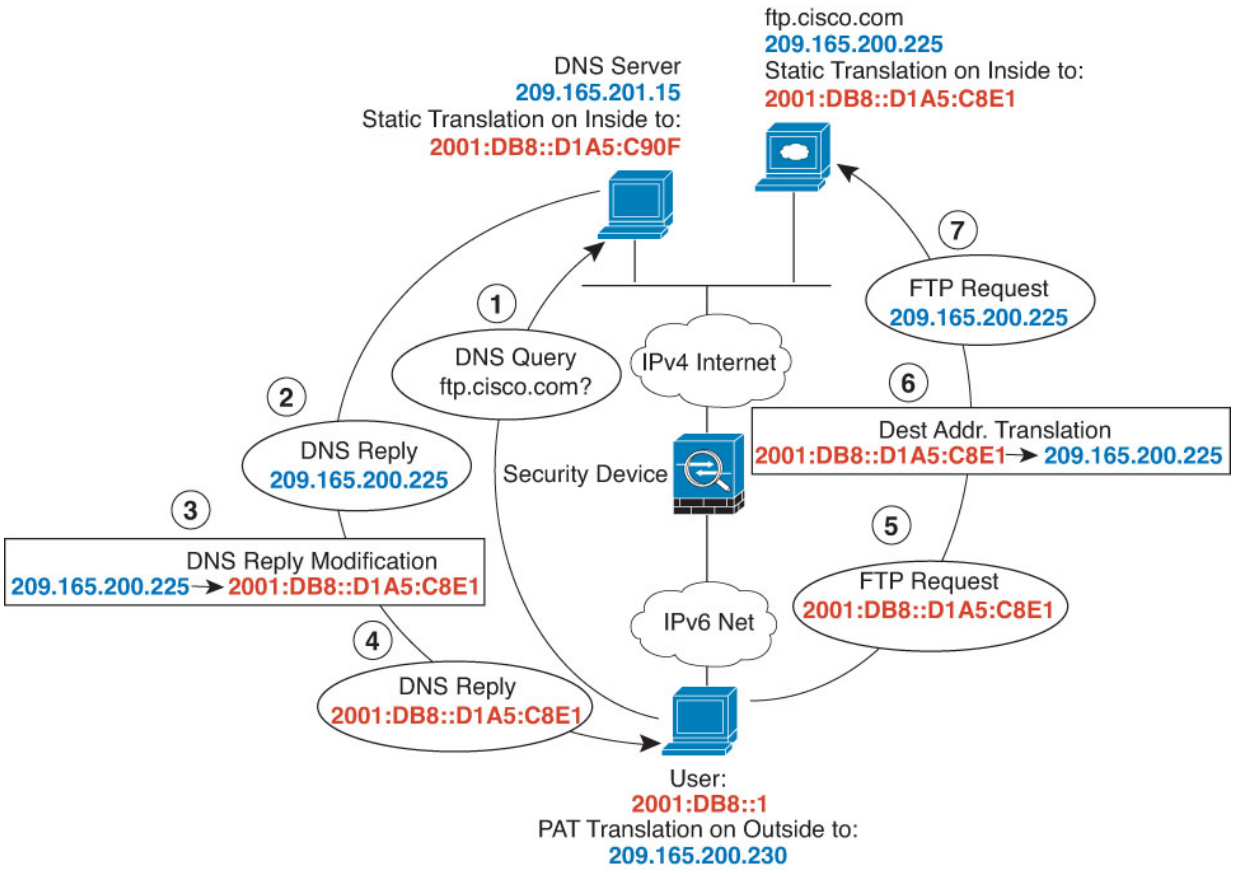

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

Step 1 Create the network objects for the FTP server, DNS server, inside network, and PAT pool.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the real FTP server address.

Name the network object (for example, ftp_server) and enter the host address, 209.165.200.225.

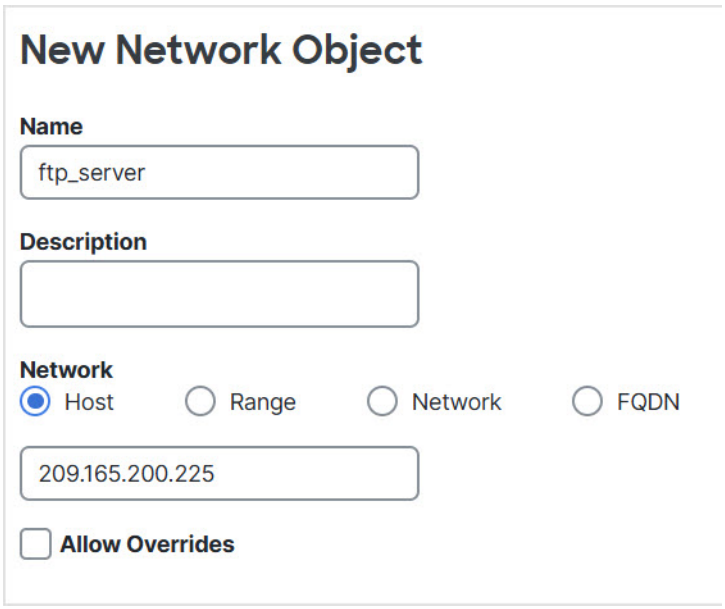

d) Click **Save**.

e) Click **Add Network** > **Add Object** and define the FTP server's translated IPv6 address.

Name the network object (for example, ftp_server_v6) and enter the host address, 2001:DB8::D1A5:C8E1.

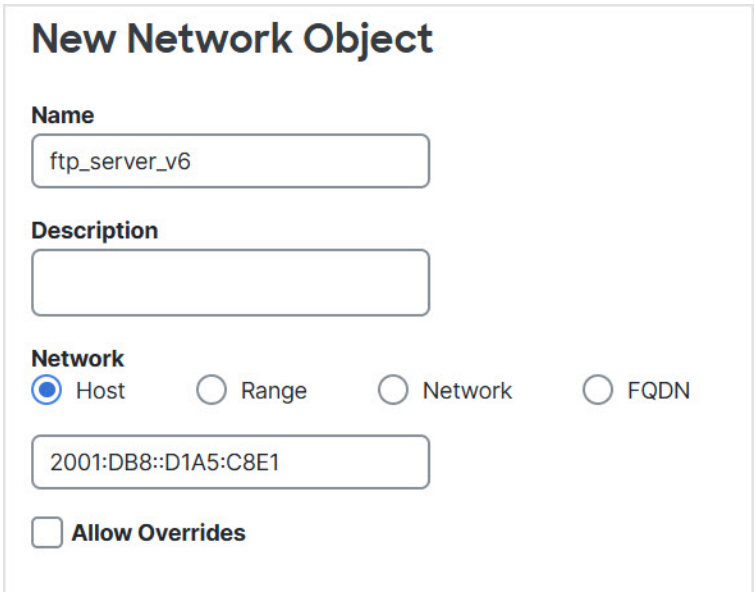

f) Click **Save**.

g) Click **Add Network** > **Add Object** and define the DNS server's real address.

Name the network object (for example, dns_server) and enter the host address, 209.165.201.15.

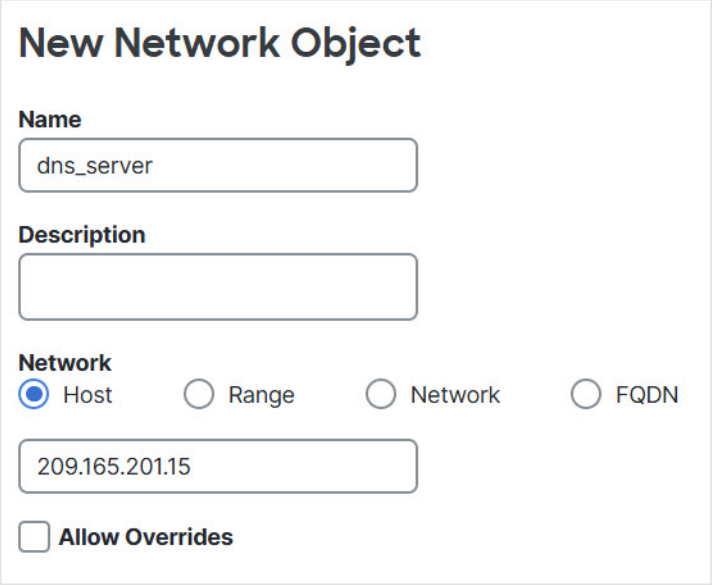

h) Click **Save**.

i) Click **Add Network** > **Add Object** and define the DNS server's translated IPv6 address.

Name the network object (for example, dns_server_v6) and enter the host address, 2001:DB8::D1A5:C90F (where D1A5:C90F is the IPv6 equivalent of 209.165.201.15).

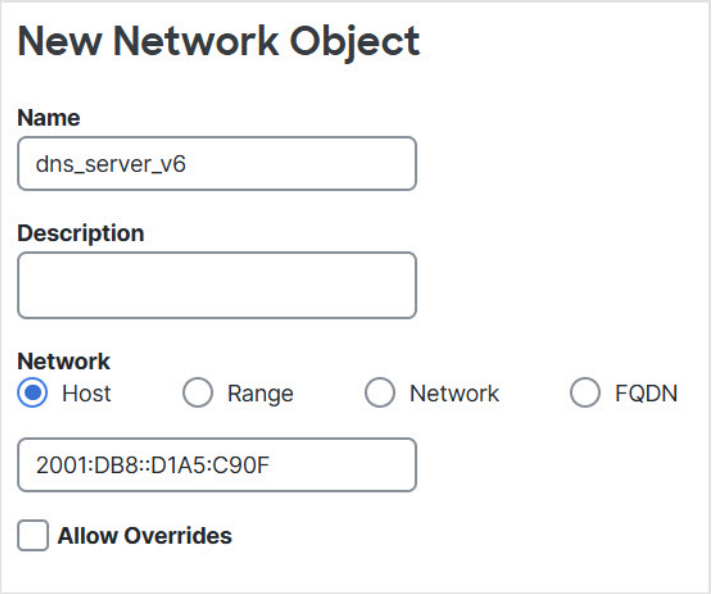

j) Click **Save**.

k) Click **Add Network** > **Add Object** and define the inside IPv6 network.

Name the network object (for example, inside_v6) and enter the network address, 2001:DB8::/96.

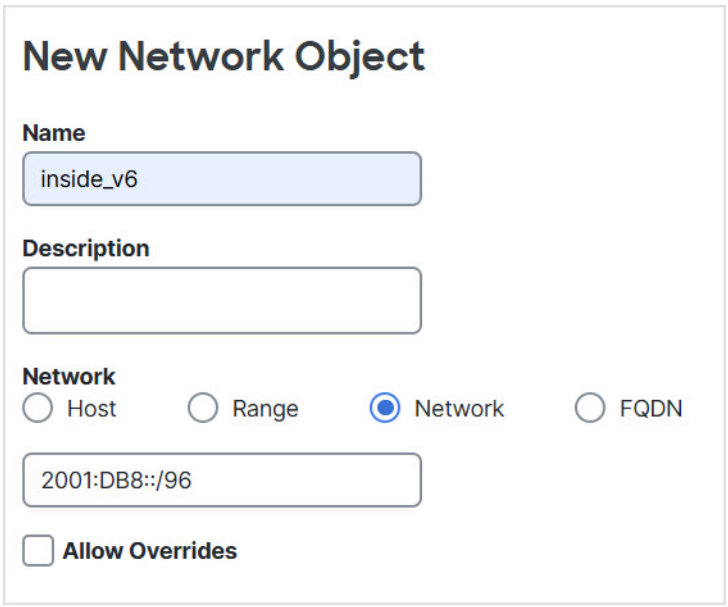

l) Click **Save**.

m) Click **Add Network** > **Add Object** and define the IPv4 PAT pool for the inside IPv6 network.

Name the network object (for example, ipv4_pool) and enter the range 209.165.200.230-209.165.200.235.

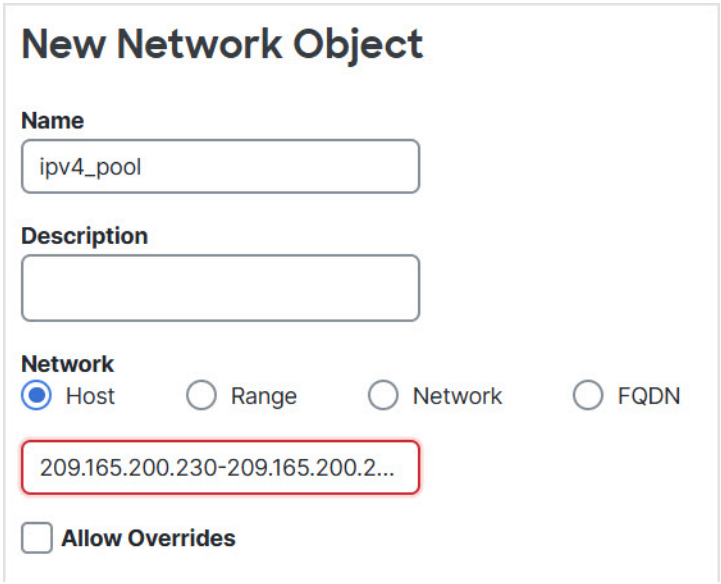

n) Click **Save**.

Step 2 Configure the static NAT rule with DNS modification for the FTP server.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.

- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.
- **Destination Interface Objects** = inside.

e) On **Translation**, configure the following:

- **Original Source** = ftp_server network object.
- **Translated Source** > **Address** = ftp_server_v6 network object.

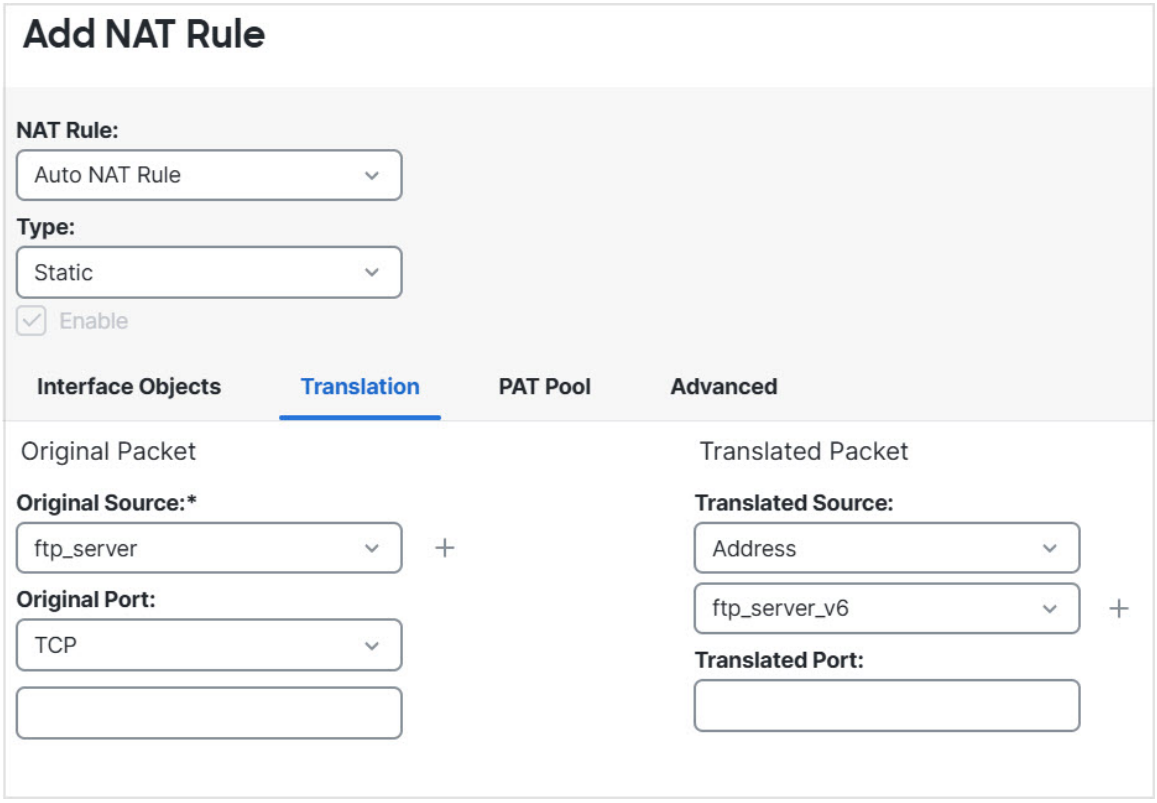

f) On **Advanced**, select the following options:

- **Translate DNS replies that match this rule**.
- **Net to Net Mapping**, because this is a one-to-one NAT46 translation.

g) Click **OK**.

Step 3 Configure the static NAT rule for the DNS server.

a) Click **Add Rule**.

b) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Static.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.
- **Destination Interface Objects** = inside.

d) On **Translation**, configure the following:

- **Original Source** = dns_server network object.
- **Translated Source** > **Address** = dns_server_v6 network object.

e) On **Advanced**, select **Net to Net Mapping**, because this is a one-to-one NAT46 translation.

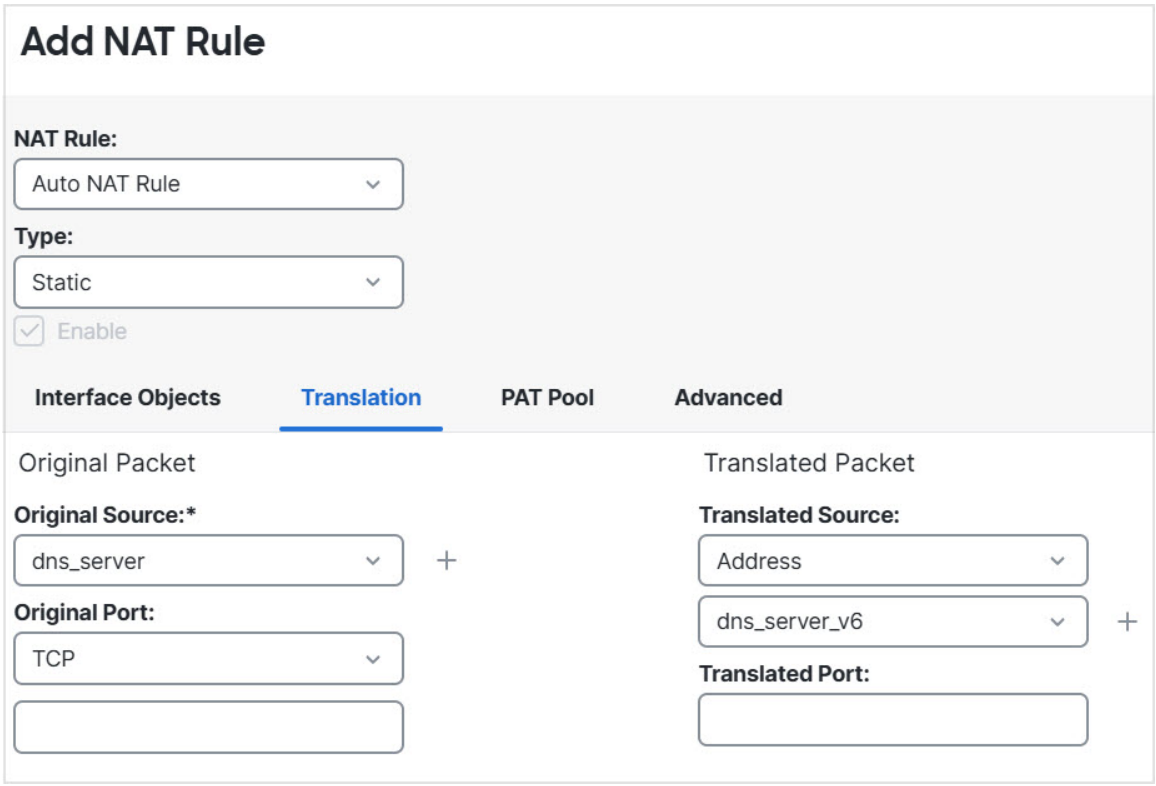

f) Click **OK**.

Step 4 Configure the dynamic NAT with a PAT pool rule for the inside IPv6 network.

a) Click **Add Rule**.

b) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.
- **Type** = Dynamic.

c) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

d) On **Translation**, configure the following:

- **Original Source** = inside_v6 network object.
- **Translated Source** > **Address** = leave this field empty.

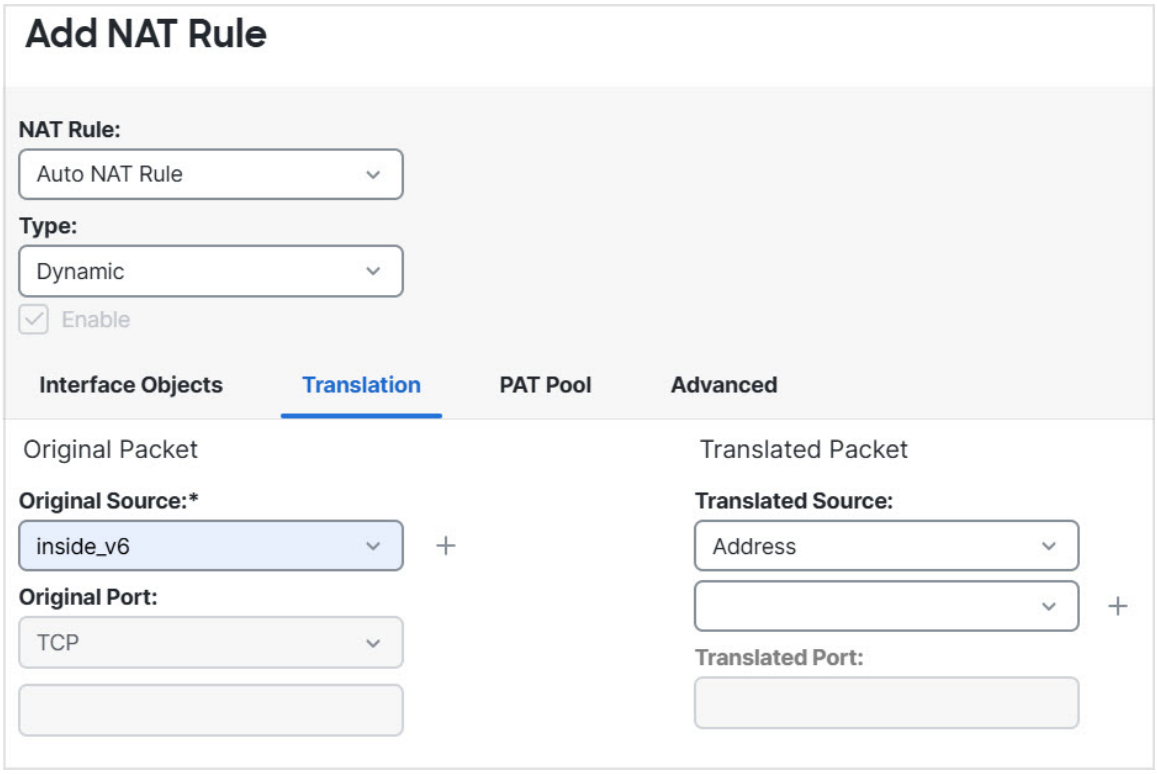

e) On **PAT Pool**, configure the following:

- **Enable PAT Pool** = select this option.
- **Translated Source** > **Address** = ipv4_pool network object.

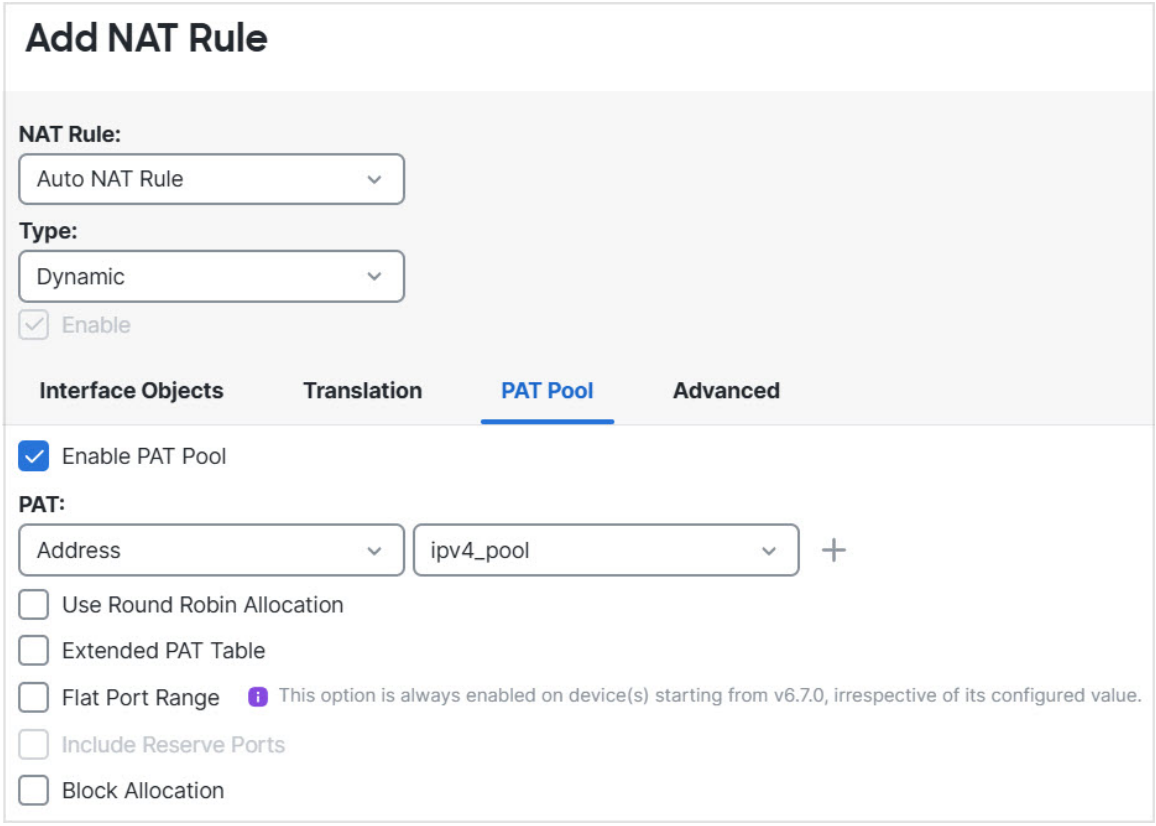

f) Click **OK**.

DNS Reply Modification, DNS Server on Outside

The following figure shows a DNS server that is accessible from the outside interface. A server, ftp.cisco.com, is on the inside interface. You configure NAT to statically translate the ftp.cisco.com real address (10.1.3.14) to a mapped address (209.165.201.10) that is visible on the outside network.

In this case, you want to enable DNS reply modification on this static rule so that inside users who have access to ftp.cisco.com using the real address receive the real address from the DNS server, and not the mapped address.

When an inside host sends a DNS request for the address of ftp.cisco.com, the DNS server replies with the mapped address (209.165.201.10). The system refers to the static rule for the inside server and translates the address inside the DNS reply to 10.1.3.14. If you do not enable DNS reply modification, then the inside host attempts to send traffic to 209.165.201.10 instead of accessing ftp.cisco.com directly.

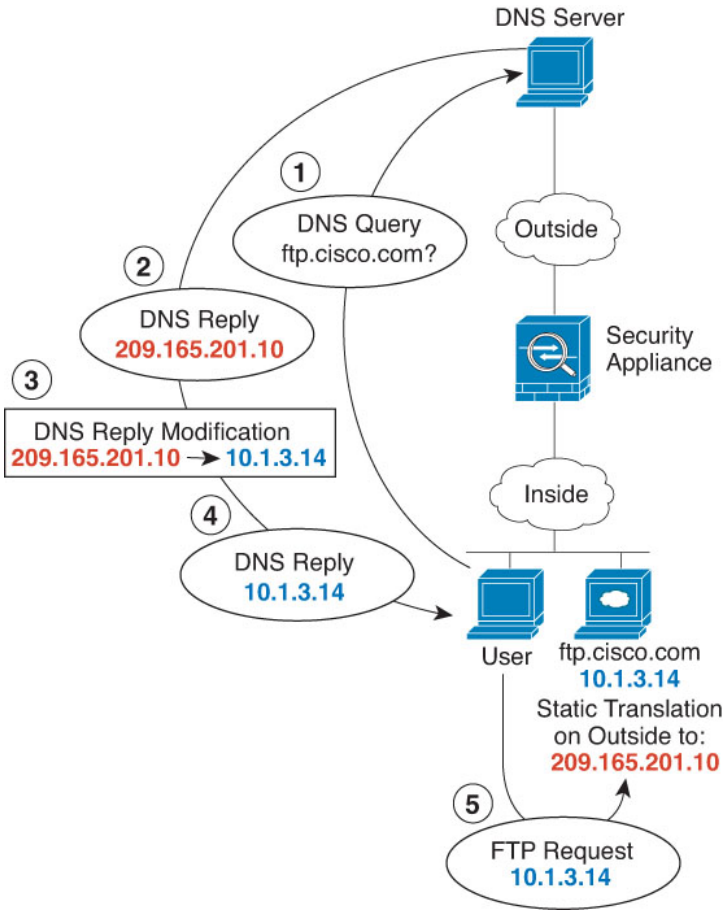

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

Step 1 Create the network objects for the FTP server.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the real FTP server address.

Name the network object (for example, ftp_server) and enter the host address, 10.1.3.14.

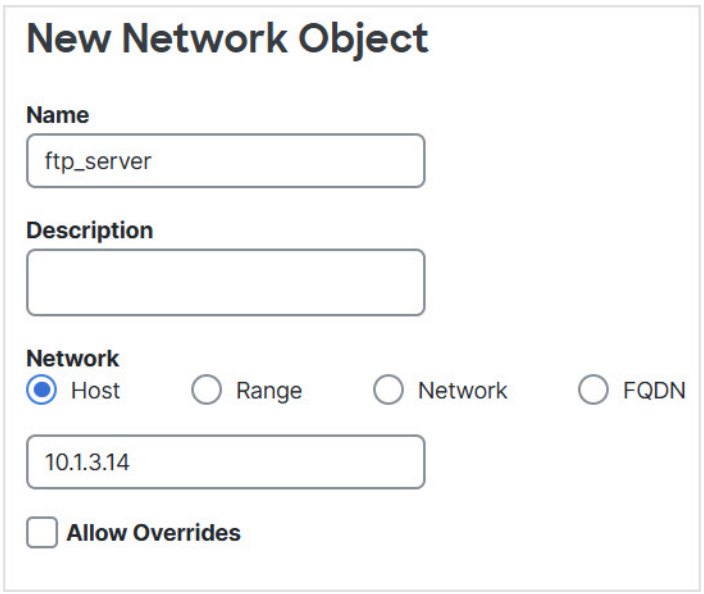

d) Click **Save**.

e) Click **Add Network** > **Add Object** and define the FTP server's translated address.

Name the network object (for example, ftp_server_outside) and enter the host address, 209.165.201.10.

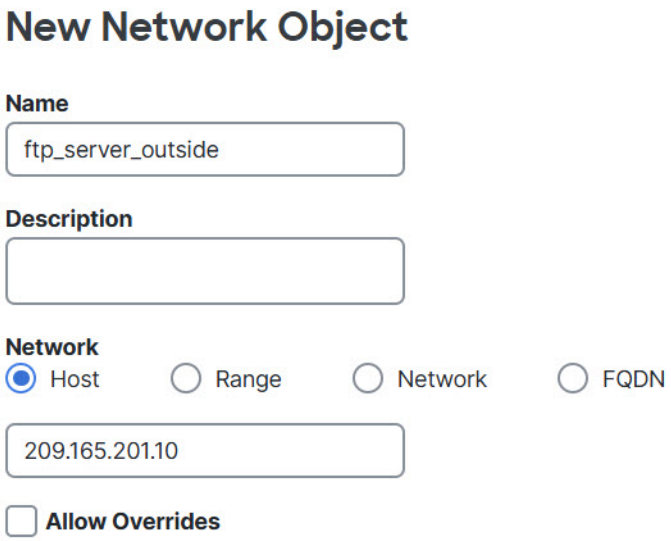

f) Click **Save**.

Step 2 Configure the static NAT rule with DNS modification for the FTP server.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.

- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = inside.
- **Destination Interface Objects** = outside.

e) On **Translation**, configure the following:

- **Original Source** = ftp_server network object.
- **Translated Source** > **Address** = ftp_server_outside network object.

f) On **Advanced**, select **Translate DNS replies that match this rule**.

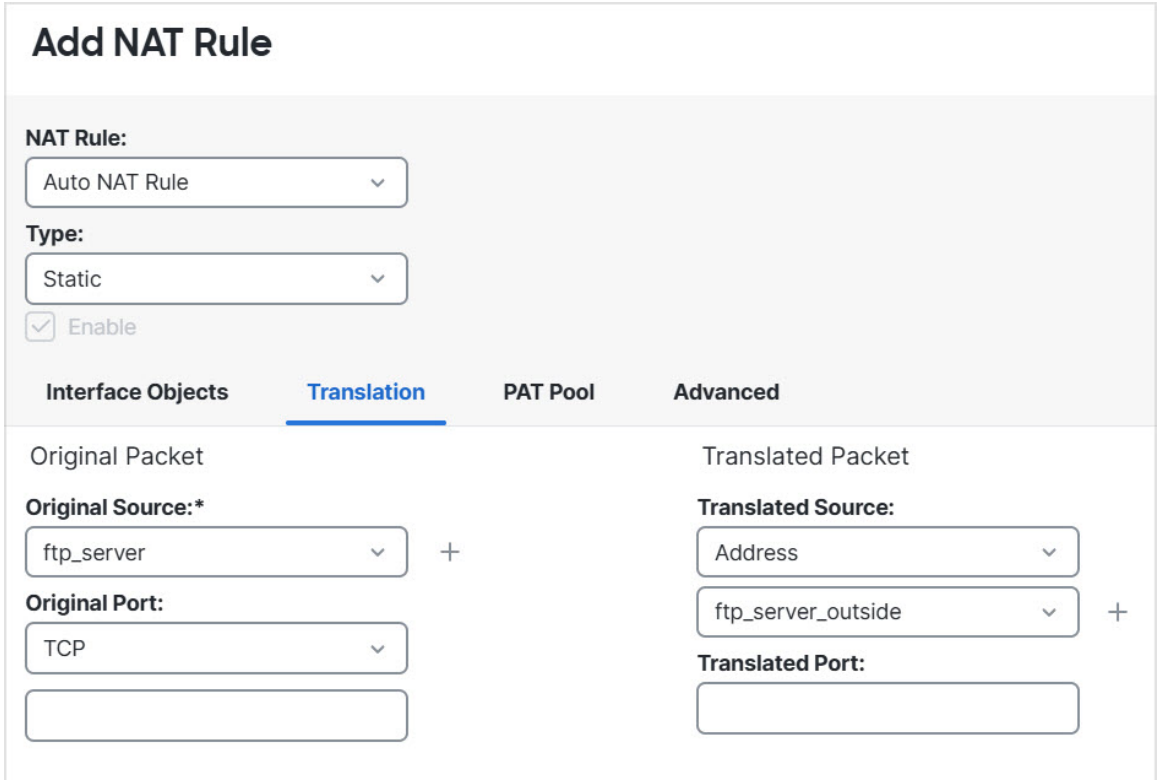

g) Click **OK**.

DNS Reply Modification, DNS Server on Host Network

The following figure shows an FTP server and DNS server on the outside. The system has a static translation for the outside server. In this case, when an inside user requests the address for ftp.cisco.com from the DNS server, the DNS server responds with the real address, 209.165.20.10. Because you want inside users to use the mapped address for ftp.cisco.com (10.1.2.56) you need to configure DNS reply modification for the static translation.

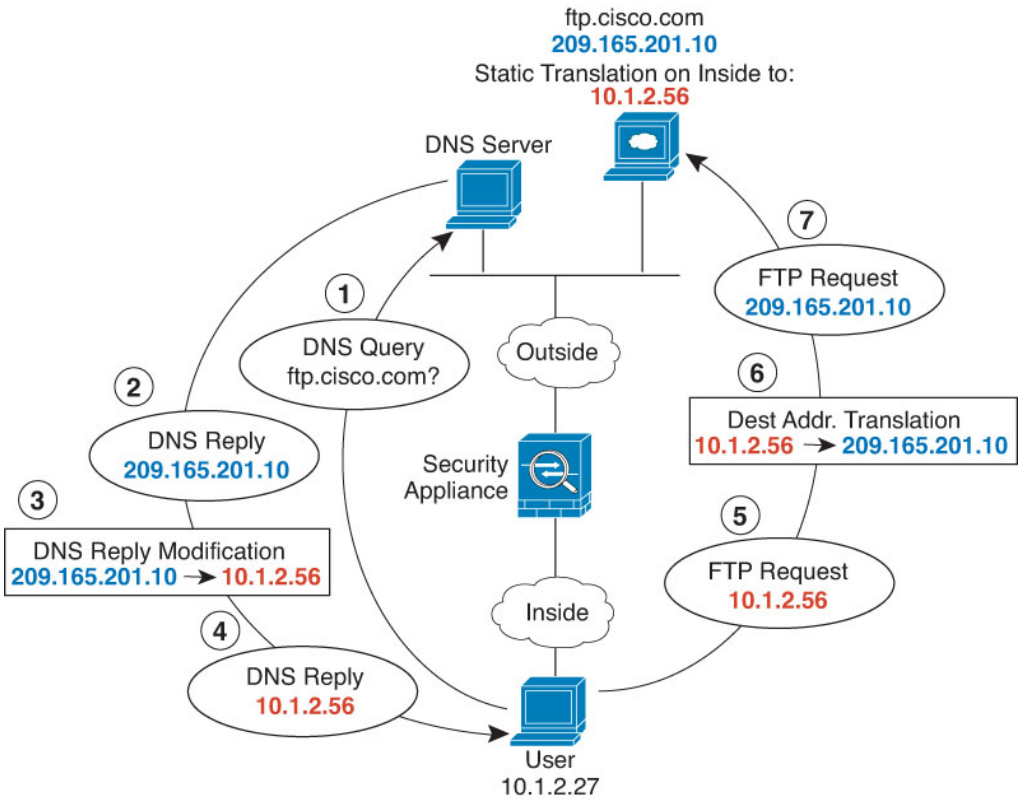

Before you begin

Ensure that you have interface objects (security zones or interface groups) that contain the interfaces for the device. In this example, we will assume the interface objects are security zones named **inside** and **outside**. To configure interface objects, select **Objects** > **Object Management**, then select **Interface**.

Procedure

Step 1 Create the network objects for the FTP server.

a) Choose **Objects** > **Object Management**.

b) Select **Network** from the table of contents and click **Add Network** > **Add Object**.

c) Define the real FTP server address.

Name the network object (for example, ftp_server) and enter the host address, 209.165.201.10.

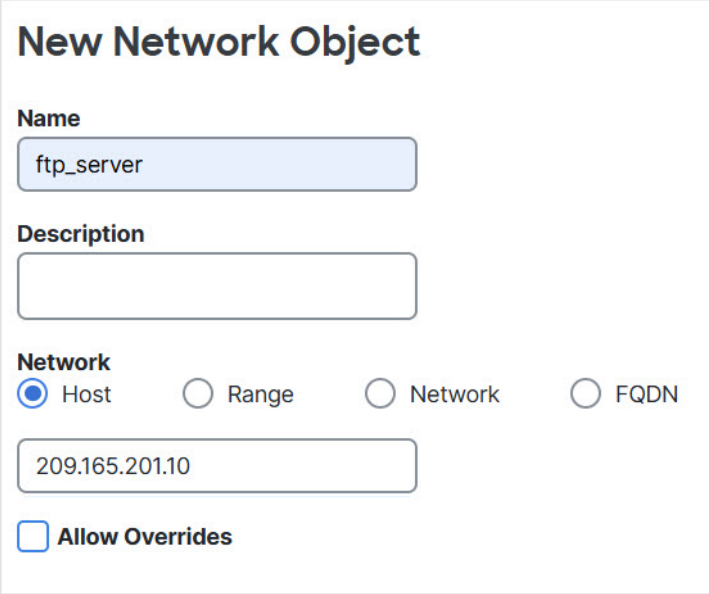

d) Click **Save**.

e) Click **Add Network** > **Add Object** and define the FTP server's translated address.

Name the network object (for example, ftp_server_translated) and enter the host address, 10.1.2.56.

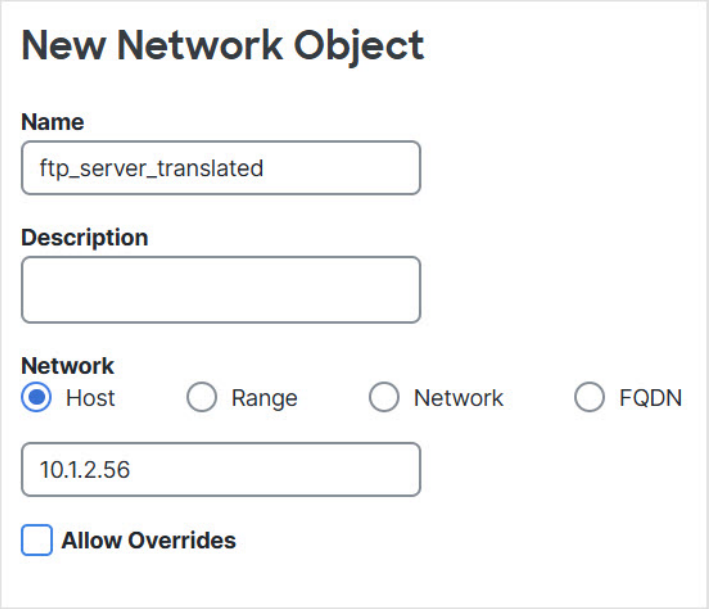

f) Click **Save**.

Step 2 Configure the static NAT rule with DNS modification for the FTP server.

a) Select **Devices** > **NAT** and create or edit the Firewall Threat Defense NAT policy.

b) Click **Add Rule**.

c) Configure the following properties:

- **NAT Rule** = Auto NAT Rule.

- **Type** = Static.

d) On **Interface Objects**, configure the following:

- **Source Interface Objects** = outside.
- **Destination Interface Objects** = inside.

e) On **Translation**, configure the following:

- **Original Source** = ftp_server network object.
- **Translated Source** > **Address** = ftp_server_translated network object.

f) On **Advanced**, select **Translate DNS replies that match this rule**.

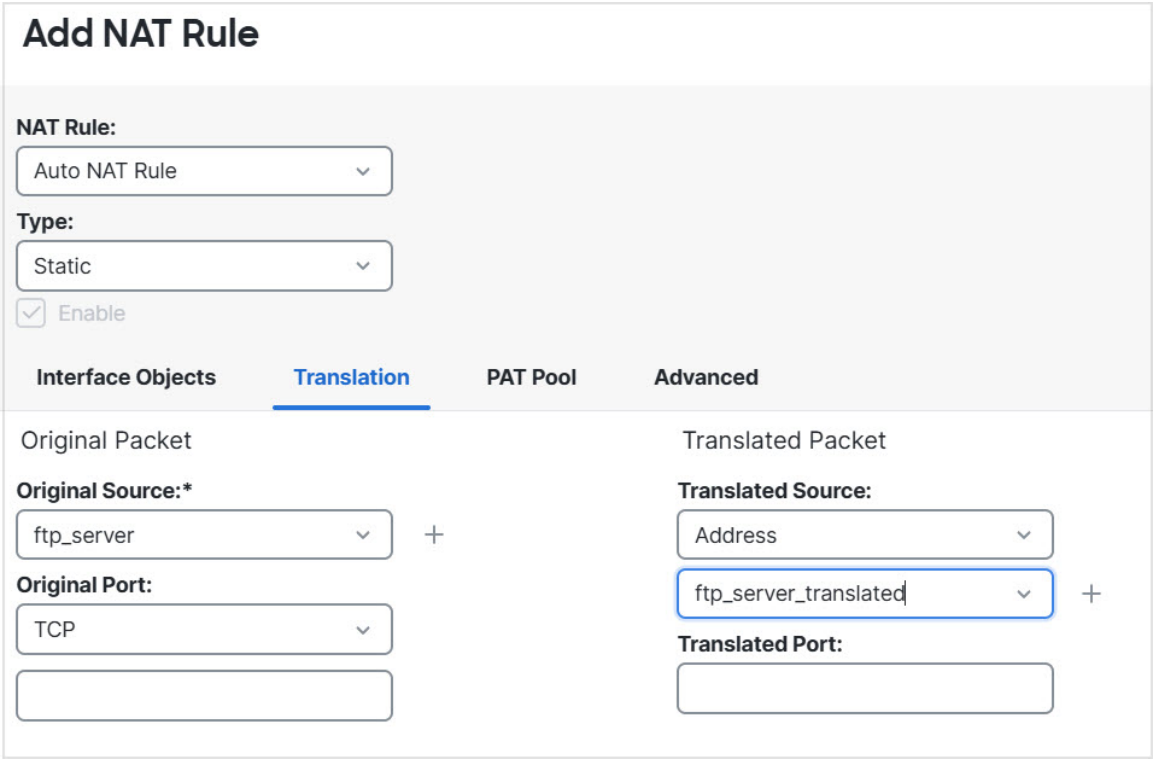

g) Click **OK**.

History for Firewall Threat Defense NAT

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Creation of network groups while editing NAT rules.	7.2.6 7.4.1	Any	You can create network groups in addition to network objects when you are editing a NAT rule. This feature is not supported in Version 7.3.x or 7.4.0.
Ability to enable, disable, or delete more than one NAT rule at a time.	7.2	Any	You can select multiple NAT rules and enable, disable, or delete them all at the same time. Enable and disable apply to manual NAT rules only, whereas delete applies to any NAT rule.
Manual NAT support for fully-qualified domain name (FQDN) objects as the translated destination.	7.1	Any	You can use an FQDN network object, such as one specifying www.example.com, as the translated destination address in manual NAT rules. The system configures the rule based on the IP address returned from the DNS server.
Changes to PAT address allocation in clustering. The PAT pool **Flat Port Range** option is now enabled by default and it is not configurable.	6.7	Any	The way PAT addresses are distributed to the members of a cluster is changed. Previously, addresses were distributed to the members of the cluster, so your PAT pool would need a minimum of one address per cluster member. Now, the control unit instead divides each PAT pool address into equal-sized port blocks and distributes them across cluster members. Each member has port blocks for the same PAT addresses. Thus, you can reduce the size of the PAT pool, even to as few as one IP address, depending on the amount of connections you typically need to PAT. Port blocks are allocated in 512-port blocks from the 1024-65535 range. You can optionally included the reserved ports, 1-1023, in this block allocation when you configure PAT pool rules. For example, in a 4-node cluster, each node gets 32 blocks with which it will be able to handle 16384 connections per PAT pool IP address compared to a single node handling all 65535 connections per PAT pool IP address. As part of this change, PAT pools for all systems, whether standalone or operating in a cluster, now use a flat port range of 1023 - 65535. Previously, you could optionally use a flat range by including the **Flat Port Range** option in a PAT pool rule. The **Flat Port Range** option is now ignored: the PAT pool is now always flat. You can optionally select the **Include Reserved Ports** option to include the 1 - 1023 port range within the PAT pool. Note that if you configure port block allocation (the **Block Allocation** PAT pool option), your block allocation size is used rather than the default 512-port block. In addition, you cannot configure extended PAT for a PAT pool for systems in a cluster.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Ability to search and filter the Firewall Threat Defense NAT rule table.	6.7	Any	You can now search for rules in Firewall Threat Defense NAT policies to help you find rules based on IP addresses, ports, object names, and so forth. Search results include partial matches. Searching on criteria filters the rule table so only matching rules are displayed. We added a search field above the rule table when you edit Firewall Threat Defense NAT policies.
Carrier grade NAT enhancements.	6.5	Any	For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time (see RFC 6888). New/modified screens: We added the **Block Allocation** option to the NAT PAT Pool tab for Firewall Threat Defense NAT rules.
Support for network range objects in NAT for Firewall Threat Defense.	6.1.0	Any	You can now use network range objects in Firewall Threat Defense NAT rules where appropriate.
Network Address Translation (NAT) for Firewall Threat Defense.	6.0.1	Any	The NAT policy for Firewall Threat Defense was added. New/modified screens: Threat Defense was added as a type of NAT policy to the **Devices** > **NAT** page.

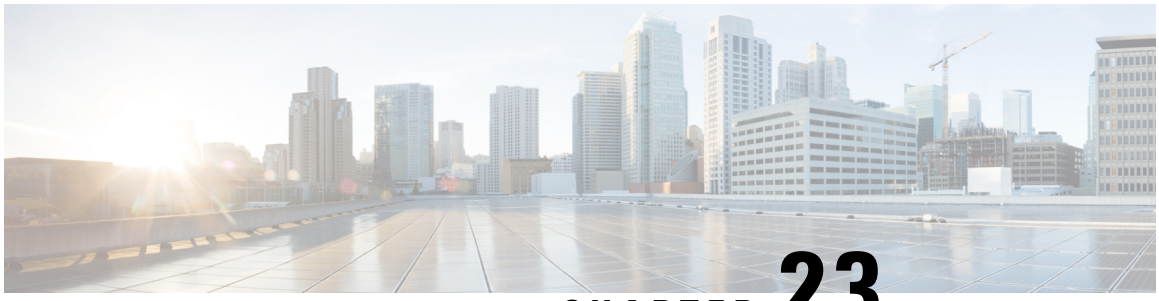

CHAPTER 23

Alarms for the Cisco ISA 3000

You can configure the alarm system on a Cisco ISA 3000 device to alert you when undesirable conditions occur.

- About Alarms, on page 1135
- Defaults for Alarms, on page 1137
- Requirements and Prerequisites for Alarms, on page 1138
- Configure the Alarms for the ISA 3000, on page 1138
- Monitoring Alarms, on page 1146
- History for Alarms, on page 1147

About Alarms

You can configure the ISA 3000 to issue alarms for a variety of conditions. If any conditions do not match the configured settings, the system triggers an alarm, which is reported by way of LEDs, syslog messages, SNMP traps, and through external devices connected to the alarm output interface. By default, triggered alarms issue syslog messages only.

You can configure the alarm system to monitor the following:

- Power supply.
- Primary and secondary temperature sensors.
- Alarm input interfaces.

The ISA 3000 has internal sensors plus two alarm input interfaces and one alarm output interface. You can connect external sensors, such as door sensors, to the alarm inputs. You can connect external alarm devices, such as buzzers or lights, to the alarm output interface.

The alarm output interface is a relay mechanism. Depending on the alarm conditions, the relay is either energized or de-energized. When it is energized, any device connected to the interface is activated. A de-energized relay results in the inactive state of any connected devices. The relay remains in an energized state as long as alarms are triggered.

For information about connecting external sensors and the alarm relay, see Cisco ISA 3000 Industrial Security Appliance Hardware Installation Guide.

Alarm Input Interfaces

You can connect the alarm input interfaces (or contacts) to external sensors, such as one that detects if a door is open.

Each alarm input interface has a corresponding LED. These LEDs convey the alarm status of each alarm input. You can configure the trigger and severity for each alarm input. In addition to the LED, you can configure the contact to trigger the output relay (to activate an external alarm), to send syslog messages, and to send SNMP traps.

The following table explains the statuses of the LEDs in response to alarm conditions for the alarm inputs. It also explains the behavior for the output relay, syslog messages, and SNMP traps, if you enable these responses to the alarm input.

Alarm Status	LED	Output Relay	Syslog	SNMP Trap
Alarm not configured	Off	—	—	—
No alarms triggered	Solid green	—	—	—
Alarm activated	Minor alarm—solid red Major alarm—flashing red	Relay energized	Syslog generated	SNMP trap sent
Alarm end	Solid green	Relay de-energized	Syslog generated	—

Alarm Output Interface

You can connect an external alarm, such as a buzzer or light, to the alarm output interface.

The alarm output interface functions as a relay and also has a corresponding LED, which conveys the alarm status of an external sensor connected to the input interface, and internal sensors such as the dual power supply and temperature sensors. You configure which alarms should activate the output relay, if any.

The following table explains the statuses of the LEDs and output relay in response to alarm conditions. It also explains the behavior for syslog messages, and SNMP traps, if you enable these responses to the alarm.

Alarm Status	LED	Output Relay	Syslog	SNMP Trap
Alarm not configured	Off	—	—	—
No alarms triggered	Solid green	—	—	—
Alarm activated	Solid red	Relay energized	Syslog generated	SNMP trap sent
Alarm end	Solid green	Relay de-energized	Syslog generated	—

Syslog Alarms

By default, the system sends syslog messages when any alarm is triggered. You can disable syslog messaging if you do not want the messages.

For syslog alarms to work, you must also enable diagnostic logging. Choose Interface**Devices** > **Platform Settings**, add or edit a Threat Defense platform settings policy that is assigned to the device, and configure destinations and settings on the **Syslog** page. For example, you can configure a syslog server, console logging, or internal buffer logging.

Without enabling a destination for diagnostic logging, the alarm system has nowhere to send syslog messages.

SNMP Alarms

You can optionally configure the alarms to send SNMP traps to your SNMP server. For SNMP trap alarms to work, you must also configure SNMP settings.

Choose **Devices** > **Platform Settings**, add or edit a Threat Defense platform settings policy that is assigned to the device, and enable SNMP and configure settings on the **SNMP** page.

Defaults for Alarms

The following table specifies the defaults for alarm input interfaces (contacts), redundant power supply, and temperature.

	Alarm	Trigger	Severity	SNMP Trap	Output Relay	Syslog Message
Alarm Contact 1	Enabled	Closed State	Minor	Disabled	Disabled	Enabled
Alarm Contact 2	Enabled	Closed State	Minor	Disabled	Disabled	Enabled
Redundant Power Supply (when enabled)	Enabled	—	—	Disabled	Disabled	Enabled
Temperature	Enabled for the primary temperature alarm (default values of 92°C and -40°C for the high and low thresholds respectively) Disabled for the secondary alarm.	—	—	Enabled for primary temperature alarm	Enabled for primary temperature alarm	Enabled for primary temperature alarm

Requirements and Prerequisites for Alarms

Model Support

Threat Defense on the ISA 3000.

Supported Domains

Any

User Roles

Admin

Configure the Alarms for the ISA 3000

You use FlexConfig to configure alarms for the ISA 3000. The following topics explain how to configure the different types of alarms.

Configure Alarm Input Contacts

If you connect the alarm input contacts (interfaces) to external sensors, you can configure the contacts to issue alarms based on the input from the sensor. In fact, the contacts are enabled by default to send syslog messages if the contact is closed, that is, if the electrical current stops flowing through the contact. You need to configure the contact only if the defaults do not meet your requirements.

The alarm contacts are numbered 1 and 2, so you need to understand how you have wired the physical pins to configure the correct settings. You configure the contacts separately.

Procedure

Step 1 Create the FlexConfig object to configure the alarm input contacts.

a) Choose **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object**.

b) Click **Add FlexConfig Object**, configure the following properties, and click **Save**.

- **Name**—The object name. For example, Configure_Alarm_Contacts.
- **Deployment**—Select **Everytime**. You want this configuration to be sent in every deployment to ensure it remains configured.
- **Type**—Keep the default, **Append**. The commands are sent to the device after the commands for directly-supported features.
- **Object body**—In the object body, type the commands needed to configure the alarm contacts. The following steps explain the commands.

c) Configure a description for the alarm contact.

alarm contact {**1** | **2**} **description** *string*

For example, to set the description of contact 1 to "Door Open," enter the following:

```
alarm contact 1 description Door Open
```

d) Configure the severity for the alarm contact.

alarm contact {**1** | **2** | **any**} **severity** {**major** | **minor** | **none**}

Instead of configuring one contact, you can specify **any** to change the severity for all contacts. The severity controls the behavior of the LED associated with the contact.

- **major**—The LED blinks red.
- **minor**—The LED is solid red. This is the default.
- **none**—The LED is off.

For example, to set the severity of contact 1 to Major, enter the following:

```
alarm contact 1 severity major
```

e) Configure the trigger for the alarm contact.

alarm contact {**1** | **2** | **any**} **trigger** {**open** | **closed**}

Instead of configuring one contact, you can specify **any** to change the trigger for all contacts. The trigger determines the electrical condition that signals an alert.

- **open**—The normal condition for the contact is closed, that is, the electrical current is running through the contact. An alert is triggered if the contact becomes open, that is, the electrical current stops flowing.
- **closed**—The normal condition for the contact is open, that is, the electrical current does not run through the contact. An alert is triggered if the contact becomes closed, that is, the electrical current starts running through the contact. This is the default.

For example, you connect a door sensor to alarm input contact 1, and its normal state has no electrical current flowing through the alarm contact (it is open). If the door is opened, the contact is closed and electrical current flows through the alarm contact. You would set the alarm trigger to closed so that the alarm goes off when the electrical current starts flowing.

```
alarm contact 1 trigger closed
```

f) Configure the actions to take when the alarm contact is triggered.

alarm facility input-alarm {**1** | **2**} {**relay** | **syslog** | **notifies**}

You can configure more than one action. For example, you can configure the device to activate the external alarm, send syslog messages, and also send SNMP traps.

- **relay**—Energize the alarm output relay, which activates the external alarm that you attached to it, such as a buzzer or a flashing light. The output LED also goes red.
- **syslog**—Send a syslog message. This option is enabled by default.
- **notifies**—Send an SNMP trap.

For example, to enable all actions for the alarm input contact 1, enter the following:

```
alarm facility input-alarm 1 relay
alarm facility input-alarm 1 syslog
alarm facility input-alarm 1 notifies
```

g) Verify that the object body contains the commands you want.

For example, if your template includes all of the command examples shown in this procedure, the object body would have the following commands:

```
alarm contact 1 description Door Open
alarm contact 1 severity major
alarm contact 1 trigger closed
alarm facility input-alarm 1 relay
alarm facility input-alarm 1 syslog
alarm facility input-alarm 1 notifies
```

The object body should look similar to the following:

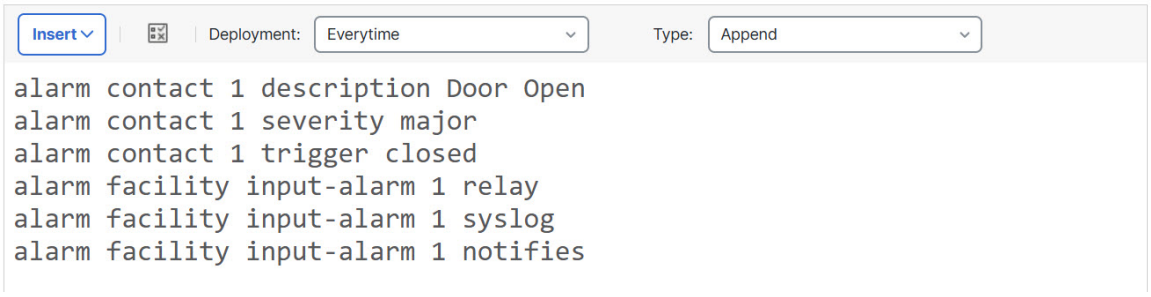

h) Click **Save**.

Step 2 Create the FlexConfig policy and assign it to the devices.

a) Choose **Devices** > **FlexConfig**.

b) Either click **New Policy**, or if an existing FlexConfig policy should be assigned to (or is already assigned to) the target devices, simply edit that policy.

When creating a new policy, assign the target devices to the policy in the dialog box where you name the policy.

c) Select the alarm contact FlexConfig object in the **User Defined** folder in the table of contents and click > to add it to the policy.

The object should be added to the **Selected Appended FlexConfigs** list.

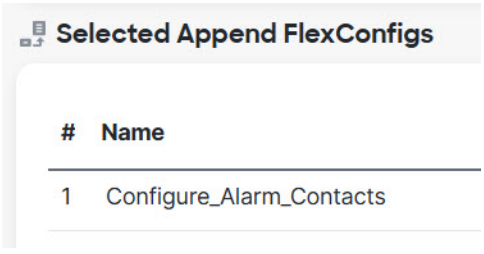

d) Click **Save**.

e) If you have not yet assigned all the targeted devices to the policy, click the **Policy Assignments** link below Save and make the assignments now.

f) Click **Preview Config**, and in the Preview dialog box, select one of the assigned devices.

The system generates a preview of the configuration CLI that will be sent to the device. Verify that the commands generated from the FlexConfig object look correct. These will be shown at the end of the preview. Note that you will also see commands generated from other changes you have made to managed features. For the alarm contact commands, you should see something similar to the following:

```
###Flex-config Appended CLI ###
alarm contact 1 description Door Open
alarm contact 1 severity major
alarm contact 1 trigger closed
alarm facility input-alarm 1 relay
alarm facility input-alarm 1 syslog
alarm facility input-alarm 1 notifies
```

Step 3 Deploy your changes.

Because you assigned a FlexConfig policy to the devices, you will always get a deployment warning, which is meant to caution you about the use of FlexConfig. Click **Proceed** to continue with the deployment.

After the deployment completes, you can check the deployment history and view the transcript for the deployment. This is especially valuable if the deployment fails. See Verify the Deployed Configuration, on page 2698.

Configure Power Supply Alarms

The ISA 3000 has two power supplies. By default, the system operates in single-power mode. However, you can configure the system to operate in dual mode, where the second power supply automatically provides power if the primary power supply fails. When you enable dual-mode, the power supply alarm is automatically enabled to send syslog alerts, but you can disable the alert altogether, or also enable SNMP traps or the alarm hardware relay.

The following procedure explains how to enable dual mode, and how to configure the power supply alarms.

Procedure

Step 1 Create the FlexConfig object to configure the power supply alarm.

a) Choose **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object**.

b) Click **Add FlexConfig Object**, configure the following properties, and click **Save**.

- **Name**—The object name. For example, Power_Supply_Alarms.
- **Deployment**—Select **Everytime**. You want this configuration to be sent in every deployment to ensure it remains configured.
- **Type**—Keep the default, **Append**. The commands are sent to the device after the commands for directly-supported features.

- **Object body**—In the object body, type the commands needed to configure the power supply alarms. The following steps explain the commands.

c) Enable dual power supply mode.

power-supply dual

For example:

```
power-supply dual
```

d) Configure the actions to take when the power supply alarm is triggered.

alarm facility power-supply rps {**relay** | **syslog** | **notifies** | **disable**}

You can configure more than one action. For example, you can configure the device to activate the external alarm, send syslog messages, and also send SNMP traps.

- **relay**—Energize the alarm output relay, which activates the external alarm that you attached to it, such as a buzzer or a flashing light. The output LED also goes red.
- **syslog**—Send a syslog message. This option is enabled by default.
- **notifies**—Send an SNMP trap.
- **disable**—Disable the power supply alarm. Any other actions configured for the power supply alarm are inoperable.

For example, to enable all actions for the power supply alarm, enter the following:

```
alarm facility power-supply rps relay
alarm facility power-supply rps syslog
alarm facility power-supply rps notifies
```

e) Verify that the object body contains the commands you want.

For example, if your template includes all of the command examples shown in this procedure, the object body would have the following commands:

```
power-supply dual
alarm facility power-supply rps relay
alarm facility power-supply rps syslog
alarm facility power-supply rps notifies
```

The object body should look similar to the following:

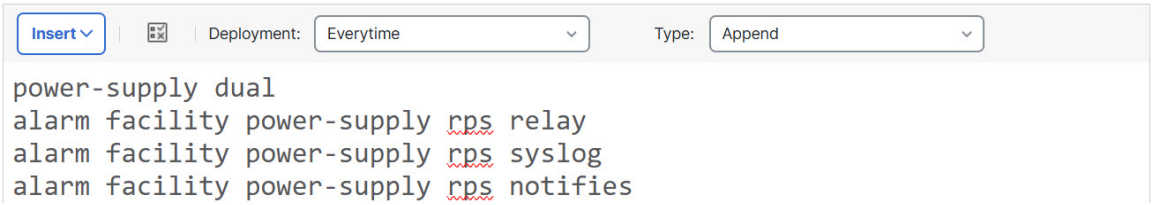

f) Click **Save**.

Step 2 Create the FlexConfig policy and assign it to the devices.

a) Choose **Devices** > **FlexConfig**.

b) Either click **New Policy**, or if an existing FlexConfig policy should be assigned to (or is already assigned to) the target devices, simply edit that policy.

When creating a new policy, assign the target devices to the policy in the dialog box where you name the policy.

c) Select the power supply alarm FlexConfig object in the **User Defined** folder in the table of contents and click > to add it to the policy.

The object should be added to the **Selected Appended FlexConfigs** list.

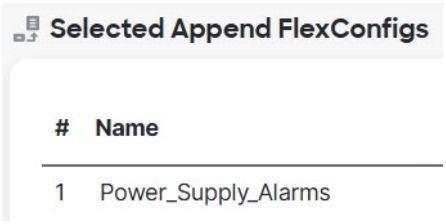

d) Click **Save**.

e) If you have not yet assigned all the targeted devices to the policy, click the **Policy Assignments** link below Save and make the assignments now.

f) Click **Preview Config**, and in the Preview dialog box, select one of the assigned devices.

The system generates a preview of the configuration CLI that will be sent to the device. Verify that the commands generated from the FlexConfig object look correct. These will be shown at the end of the preview. Note that you will also see commands generated from other changes you have made to managed features. For the power supply alarm commands, you should see something similar to the following:

```
###Flex-config Appended CLI ###
power-supply dual
alarm facility power-supply rps relay
alarm facility power-supply rps syslog
alarm facility power-supply rps notifies
```

Step 3 Deploy your changes.

Because you assigned a FlexConfig policy to the devices, you will always get a deployment warning, which is meant to caution you about the use of FlexConfig. Click **Proceed** to continue with the deployment.

After the deployment completes, you can check the deployment history and view the transcript for the deployment. This is especially valuable if the deployment fails. See Verify the Deployed Configuration, on page 2698.

Configure Temperature Alarms

You can configure alarms based on the temperature of the CPU card in the device.

You can set a primary and secondary temperature range. If the temperature drops below the low threshold, or exceeds the high threshold, the alarm is triggered.

The primary temperature alarm is enabled by default for all alarm actions: output relay, syslog, and SNMP. The default settings for the primary temperature range is -40°C to 92°C.

The secondary temperature alarm is disabled by default. You can set the secondary temperature within the range -35°C to 85°C.

Because the secondary temperature range is more restrictive than the primary range, if you set either the secondary low or high temperature, that setting disables the corresponding primary setting, even if you configure non-default values for the primary setting. You cannot enable two separate high and two separate low temperature alarms.

Thus, in practice, you should configure the primary only, or the secondary only, setting for high and low.

Procedure

Step 1 Create the FlexConfig object to configure the temperature alarms.

a) Choose **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object**.

b) Click **Add FlexConfig Object**, configure the following properties, and click **Save**.

- **Name**—The object name. For example, Configure_Temperature_Alarms.
- **Deployment**—Select **Everytime**. You want this configuration to be sent in every deployment to ensure it remains configured.
- **Type**—Keep the default, **Append**. The commands are sent to the device after the commands for directly-supported features.
- **Object body**—In the object body, type the commands needed to configure the temperature alarms. The following steps explain the commands.

c) Configure the acceptable temperature range.

alarm facility temperature {**primary** | **secondary**} {**low** | **high**} *temperature*

The temperature is in Celsius. The allowed range for the primary alarm is -40 to 92, which is also the default range. The allowed range for the secondary alarm is -35 to 85. The low value must be lower than the high value.

For example, to set a more restrictive temperature range of -20 to 80, which falls within the allowed range for the secondary alarm, configure the secondary alarm as follows:

```
alarm facility temperature secondary low -20
alarm facility temperature secondary high 80
```

d) Configure the actions to take when the temperature alarm is triggered.

alarm facility temperature {**primary** | **secondary**} {**relay** | **syslog** | **notifies**}

You can configure more than one action. For example, you can configure the device to activate the external alarm, send syslog messages, and also send SNMP traps.

- **relay**—Energize the alarm output relay, which activates the external alarm that you attached to it, such as a buzzer or a flashing light. The output LED also goes red.
- **syslog**—Send a syslog message.
- **notifies**—Send an SNMP trap.

For example, to enable all actions for the secondary temperature alarm, enter the following:

```
alarm facility temperature secondary relay
alarm facility temperature secondary syslog
alarm facility temperature secondary notifies
```

e) Verify that the object body contains the commands you want.

For example, if your template includes all of the command examples shown in this procedure, the object body would have the following commands:

```
alarm facility temperature secondary low -20
alarm facility temperature secondary high 80
alarm facility temperature secondary relay
alarm facility temperature secondary syslog
alarm facility temperature secondary notifies
```

The object body should look similar to the following:

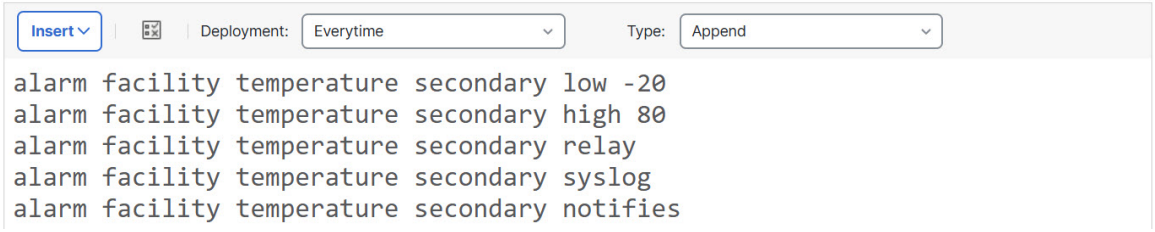

f) Click **Save**.

Step 2 Create the FlexConfig policy and assign it to the devices.

a) Choose **Devices** > **FlexConfig**.

b) Either click **New Policy**, or if an existing FlexConfig policy should be assigned to (or is already assigned to) the target devices, simply edit that policy.

When creating a new policy, assign the target devices to the policy in the dialog box where you name the policy.

c) Select the temperature alarms FlexConfig object in the **User Defined** folder in the table of contents and click > to add it to the policy.

The object should be added to the **Selected Appended FlexConfigs** list.

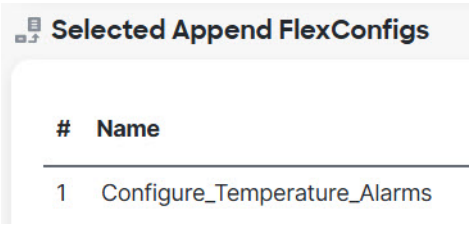

d) Click **Save**.

e) If you have not yet assigned all the targeted devices to the policy, click the **Policy Assignments** link below Save and make the assignments now.

f) Click **Preview Config**, and in the Preview dialog box, select one of the assigned devices.

The system generates a preview of the configuration CLI that will be sent to the device. Verify that the commands generated from the FlexConfig object look correct. These will be shown at the end of the

preview. Note that you will also see commands generated from other changes you have made to managed features. For the temperature alarms commands, you should see something similar to the following:

```
###Flex-config Appended CLI ###
alarm facility temperature secondary low -20
alarm facility temperature secondary high 80
alarm facility temperature secondary relay
alarm facility temperature secondary syslog
alarm facility temperature secondary notifies
```

Step 3 Deploy your changes.

Because you assigned a FlexConfig policy to the devices, you will always get a deployment warning, which is meant to caution you about the use of FlexConfig. Click **Proceed** to continue with the deployment.

After the deployment completes, you can check the deployment history and view the transcript for the deployment. This is especially valuable if the deployment fails. See Verify the Deployed Configuration, on page 2698.

Monitoring Alarms

The following topics explain how to monitor and manage alarms.

Monitoring Alarm Status

You can use the following commands in the CLI to monitor alarms.

- **show alarm settings**

  Shows the current configuration for each possible alarm.

- **show environment alarm-contact**

  Shows information about the physical status of the input alarm contacts.

- **show facility-alarm relay**

  Shows information about the alarms that have triggered the output relay.

- **show facility-alarm status** [**info** | **major** | **minor**]

  Shows information on all alarms that have been triggered. You can limit the view by filtering on **major** or **minor** status. The **info** keyword provides the same output as using no keyword.

Monitoring Syslog Messages for Alarms

Depending on the type of alarms you configure, you might see the following syslog messages.

Dual Power Supply Alarms

- %FTD-1-735005: Power Supply Unit Redundancy OK
- %FTD-1-735006: Power Supply Unit Redundancy Lost

Temperature Alarms

In these alarms, *Celsius* is replaced by the temperature detected on the device, in Celsius.

- %FTD-6-806001: Primary alarm CPU temperature is High *Celsius*
- %FTD-6-806002: Primary alarm for CPU high temperature is cleared
- %FTD-6-806003: Primary alarm CPU temperature is Low *Celsius*
- %FTD-6-806004: Primary alarm for CPU Low temperature is cleared
- %FTD-6-806005: Secondary alarm CPU temperature is High *Celsius*
- %FTD-6-806006: Secondary alarm for CPU high temperature is cleared
- %FTD-6-806007: Secondary alarm CPU temperature is Low *Celsius*
- %FTD-6-806008: Secondary alarm for CPU Low temperature is cleared

Alarm Input Contact Alarms

In these alarms, *description* is the description for the contact that you configured.

- %FTD-6-806009: Alarm asserted for ALARM_IN_1 *alarm_1_description*
- %FTD-6-806010: Alarm cleared for ALARM_IN_1 *alarm_1_description*
- %FTD-6-806011: Alarm asserted for ALARM_IN_2 *alarm_2_description*
- %FTD-6-806012: Alarm cleared for ALARM_IN_2 *alarm_2_description*

Turning Off the External Alarm

If you are using an external alarm that is attached to the alarm output, and the alarm is triggered, you can turn off the external alarm from the device CLI using the **clear facility-alarm output** command. This command de-energizes the output pin and also turns off the output LED.

History for Alarms

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Description
Alarms for the Cisco ISA 3000 series.	6.7	Any	Configuring alarms for the Cisco ISA 3000 series was validated using FlexConfig. You should be able to configure the alarms in older releases that support FlexConfig, except for the dual power supply alarms. Supported platforms: Secure Firewall Threat Defense on the ISA 3000.

PART IV

Routing

- Static and Default Routes, on page 1151
- Virtual Routers, on page 1167
- ECMP, on page 1223
- Bidirectional Forwarding Detection Routing, on page 1233
- OSPF, on page 1239
- EIGRP, on page 1267
- BGP, on page 1277
- RIP, on page 1297
- Multicast, on page 1303
- Policy Based Routing, on page 1323

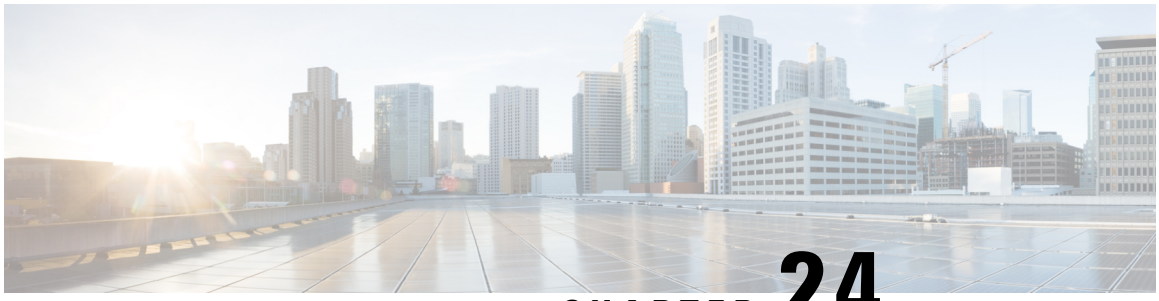

CHAPTER 24

Static and Default Routes

This chapter describes how to configure static and default routes on the Firewall Threat Defense.

- About Static and Default Routes, on page 1151
- Requirements and Prerequisites for Static Routes, on page 1153
- Guidelines for Static and Default Routes, on page 1154
- Add a Static Route, on page 1154
- Reference for Routing, on page 1156

About Static and Default Routes

To route traffic to a non-connected host or network, you must define a route to the host or network, either using static or dynamic routing. Generally, you must configure at least one static route: a default route for all traffic that is not routed by other means to a default network gateway, typically the next hop router.

Default Route

The simplest option is to configure a default static route to send all traffic to an upstream router, relying on the router to route the traffic for you. A default route identifies the gateway IP address to which the Firewall Threat Defense device sends all IP packets for which it does not have a learned or static route. A default static route is simply a static route with 0.0.0.0/0 (IPv4) or ::/0 (IPv6) as the destination IP address.

You should always define a default route.

The Firewall Threat Defense has separate routing tables for data interfaces and for management-only interfaces (including the special Linux Management interface). You can only add a default route for the data routing table. The Firewall Threat Defense automatically adds a default route in the management-only routing table that sends traffic to the Linux Management interface, where a separate route lookup occurs in the Linux routing table. You can add static routes to the Linux routing table that can be used by Management using the Firewall Threat Defense CLI **configure network static-routes** command.

Note The *default* Linux route is set with the **configure network ipv4** or **configure network ipv6** command.

Static Routes

You might want to use static routes in the following cases:

- Your networks use an unsupported router discovery protocol.
- Your network is small and you can easily manage static routes.
- You do not want the traffic or CPU overhead associated with routing protocols.
- In some cases, a default route is not enough. The default gateway might not be able to reach the destination network, so you must also configure more specific static routes. For example, if the default gateway is outside, then the default route cannot direct traffic to any inside networks that are not directly connected to the Firewall Threat Defense device.
- You are using a feature that does not support dynamic routing protocols.
- Virtual routers use static routes to create route leaks. Route leaks enable flow of traffic from an interface of a virtual router to another interface in another virtual router. For more information, see Interconnecting Virtual Routers, on page 1171.

Route to null0 Interface to Drop Unwanted Traffic

Access rules let you filter packets based on the information contained in their headers. A static route to the null0 interface is a complementary solution to access rules. You can use a null0 route to forward unwanted or undesirable traffic so the traffic is dropped.

Static null0 routes have a favorable performance profile. You can also use static null0 routes to prevent routing loops. BGP can leverage the static null0 route for Remotely Triggered Black Hole routing.

Route Priorities

- Routes that identify a specific destination take precedence over the default route.
- When multiple routes exist to the same destination (either static or dynamic), then the administrative distance for the route determines priority. Static routes are set to 1, so they typically are the highest priority routes.
- When you have multiple static routes to the same destination with the same administrative distance, see Equal-Cost Multi-Path (ECMP) Routing, on page 1164.
- For traffic emerging from a tunnel with the Tunneled option, this route overrides any other configured or learned default routes.

Transparent Firewall Mode and Bridge Group Routes

For traffic that originates on the Firewall Threat Defense device and is destined through a bridge group member interface for a non-directly connected network, you need to configure either a default route or static routes so the Firewall Threat Defense device knows out of which bridge group member interface to send traffic. Traffic that originates on the Firewall Threat Defense device might include communications to a syslog server or SNMP server. If you have servers that cannot all be reached through a single default route, then you must configure static routes. For transparent mode, you cannot specify the BVI as the gateway interface; only

member interfaces can be used. For bridge groups in routed mode, you must specify the BVI in a static route; you cannot specify a member interface. See MAC Address vs. Route Lookups, on page 290 for more information.

Static Route Tracking

One of the problems with static routes is that there is no inherent mechanism for determining if the route is up or down. They remain in the routing table even if the next hop gateway becomes unavailable. Static routes are only removed from the routing table if the associated interface on the Firewall Threat Defense device goes down.

The static route tracking feature provides a method for tracking the availability of a static route and installing a backup route if the primary route should fail. For example, you can define a default route to an ISP gateway and a backup default route to a secondary ISP in case the primary ISP becomes unavailable.

The Firewall Threat Defense device implements static route tracking by associating a static route with a monitoring target host on the destination network that the Firewall Threat Defense device monitors using ICMP echo requests. If an echo reply is not received within a specified time period, the host is considered down, and the associated route is removed from the routing table. An untracked backup route with a higher metric is used in place of the removed route.

When selecting a monitoring target, you need to make sure that it can respond to ICMP echo requests. The target can be any network object that you choose, but you should consider using the following:

- The ISP gateway (for dual ISP support) address
- The next hop gateway address (if you are concerned about the availability of the gateway)
- A server on the target network, such as a syslog server, that the Firewall Threat Defense device needs to communicate with
- A persistent network object on the destination network

Note A PC that may be shut down at night is not a good choice.

You can configure static route tracking for statically defined routes or default routes obtained through DHCP or PPPoE. You can only enable PPPoE clients on multiple interfaces with route tracking configured.

Requirements and Prerequisites for Static Routes

Model Support

Threat Defense

Supported Domains

Any

User Roles

Admin

Network Admin

Guidelines for Static and Default Routes

Firewall Mode and Bridge Groups

- In transparent mode, static routes must use the bridge group member interface as the gateway; you cannot specify the BVI.
- In routed mode, you must specify the BVI as the gateway; you cannot specify the member interface.
- Static route tracking is not supported for bridge group member interfaces or on the BVI.

Supported Network Address

- Static route tracking is not supported for IPv6.
- The Firewall Threat Defense does not support Class E routing, so a Class E network is not routable in static routes.

Clustering

- In clustering, static route tracking is only supported on the control node.

Network Object Group

You cannot use a range of network objects or a network object group having a range of IP addresses in a static route.

ASP and RIB Route Entries

All routes and its distance installed on the device are captured in the ASP routing table. This is common for all static and dynamic routing protocols. Only the best distance route is captured in the RIB table.

Add a Static Route

A static route defines where to send traffic for specific destination networks. You should at a minimum define a default route. A default route is simply a static route with 0.0.0.0/0 as the destination IP address.

To configure routes for redundant manager access data interfaces, see Configure a Redundant Manager Access Data Interface, on page 153.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 (May be required) From the virtual routers drop-down list, select the virtual router for which you are configuring a static route.

Step 4 Select **Static Route**.

Step 5 Click **Add Routes**.

Step 6 Click **IPv4** or **IPv6** depending on the type of static route that you are adding.

Step 7 Choose the **Interface** to which this static route applies.

For transparent mode, choose a bridge group member interface name. For routed mode with bridge groups, you can choose either the bridge group member interface for the BVI name. To "black hole" unwanted traffic, choose the **Null0** interface.

For a device using virtual routing, you can select an interface that belongs to another virtual router. You can create such a static route if you want to leak traffic from this virtual router into the other virtual router. For more information, see Interconnecting Virtual Routers, on page 1171.

Step 8 In the **Available Network** list, choose the destination network.

To define a default route, create an object with the address 0.0.0.0/0 and select it here.

Note
Although you can create and choose a Network Object Group containing a range of IP addresses, the Firewall Management Center does not support using ranges in static routes.

Step 9 In the **Gateway** or **IPv6 Gateway** field, enter or choose the gateway router that is the next hop for this route. You can provide an IP address or a Networks/Hosts object.

When you are using static-route configuration for virtual routers to leak routes, do not specify the next hop gateway.

Step 10 In the **Metric** field, enter the number of hops to the destination network. Valid values range from 1 to 255; the default value is 1.

The metric is a measurement of the "expense" of a route, based on the number of hops (hop count) to the network on which a specific host resides. Hop count is the number of networks that a network packet must traverse, including the destination network, before it reaches its final destination. The metric is used to compare routes among different routing protocols. The default administrative distance for static routes is 1, giving it precedence over routes discovered by dynamic routing protocols but not directly connected routes. The default administrative distance for routes discovered by OSPF is 110. If a static route has the same administrative distance as a dynamic route, the static route takes precedence. Connected routes always take precedence over static or dynamically discovered routes.

Step 11 (Optional) For a default route, click the **Tunneled** checkbox to define a separate default route for VPN traffic.

You can define a separate default route for VPN traffic if you want your VPN traffic to use a different default route than your non VPN traffic. For example, traffic incoming from VPN connections can be easily directed towards internal networks, while traffic from internal networks can be directed towards the outside. When you create a default route with the tunneled option, all traffic from a tunnel terminating on the device that cannot be routed using learned or static routes, is sent to this route. You can configure only one default tunneled gateway per device. ECMP for tunneled traffic is not supported.

Step 12 (IPv4 static route only) To monitor route availability, enter or choose the name of an SLA (service level agreement) Monitor object that defines the monitoring policy, in the **Route Tracking** field.

See SLA Monitor, on page 1441.

Step 13 Click **Ok**.

Reference for Routing

This section describes underlying concepts of how routing behaves within the Firewall Threat Defense.

Path Determination

Routing protocols use metrics to evaluate what path will be the best for a packet to travel. A metric is a standard of measurement, such as path bandwidth, that is used by routing algorithms to determine the optimal path to a destination. To aid the process of path determination, routing algorithms initialize and maintain routing tables, which include route information. Route information varies depending on the routing algorithm used.

Routing algorithms fill routing tables with a variety of information. Destination or next hop associations tell a router that a particular destination can be reached optimally by sending the packet to a particular router representing the next hop on the way to the final destination. When a router receives an incoming packet, it checks the destination address and attempts to associate this address with a next hop.

Routing tables also can include other information, such as data about the desirability of a path. Routers compare metrics to determine optimal routes, and these metrics differ depending on the design of the routing algorithm used.

Routers communicate with one another and maintain their routing tables through the transmission of a variety of messages. The routing update message is one such message that generally consists of all or a portion of a routing table. By analyzing routing updates from all other routers, a router can build a detailed picture of network topology. A link-state advertisement, another example of a message sent between routers, informs other routers of the state of the sender links. Link information also can be used to build a complete picture of network topology to enable routers to determine optimal routes to network destinations.

Supported Route Types

There are several route types that a router can use. The Firewall Threat Defense device uses the following route types:

- Static Versus Dynamic
- Single-Path Versus Multipath
- Flat Versus Hierarchical
- Link-State Versus Distance Vector

Static Versus Dynamic

Static routing algorithms are actually table mappings established by the network administrator. These mappings do not change unless the network administrator alters them. Algorithms that use static routes are simple to design and work well in environments where network traffic is relatively predictable and where network design is relatively simple.

Because static routing systems cannot react to network changes, they generally are considered unsuitable for large, constantly changing networks. Most of the dominant routing algorithms are dynamic routing algorithms, which adjust to changing network circumstances by analyzing incoming routing update messages. If the message indicates that a network change has occurred, the routing software recalculates routes and sends out new routing update messages. These messages permeate the network, stimulating routers to rerun their algorithms and change their routing tables accordingly.

Dynamic routing algorithms can be supplemented with static routes where appropriate. A router of last resort (a default route for a router to which all unrouteable packets are sent), for example, can be designated to act as a repository for all unrouteable packets, ensuring that all messages are at least handled in some way.

Single-Path Versus Multipath

Some sophisticated routing protocols support multiple paths to the same destination. Unlike single-path algorithms, these multipath algorithms permit traffic multiplexing over multiple lines. The advantages of multipath algorithms are substantially better throughput and reliability, which is generally called load sharing.

Flat Versus Hierarchical

Some routing algorithms operate in a flat space, while others use routing hierarchies. In a flat routing system, the routers are peers of all others. In a hierarchical routing system, some routers form what amounts to a routing backbone. Packets from non-backbone routers travel to the backbone routers, where they are sent through the backbone until they reach the general area of the destination. At this point, they travel from the last backbone router through one or more non-backbone routers to the final destination.

Routing systems often designate logical groups of nodes, called domains, autonomous systems, or areas. In hierarchical systems, some routers in a domain can communicate with routers in other domains, while others can communicate only with routers within their domain. In very large networks, additional hierarchical levels may exist, with routers at the highest hierarchical level forming the routing backbone.

The primary advantage of hierarchical routing is that it mimics the organization of most companies and therefore supports their traffic patterns well. Most network communication occurs within small company groups (domains). Because intradomain routers need to know only about other routers within their domain, their routing algorithms can be simplified, and, depending on the routing algorithm being used, routing update traffic can be reduced accordingly.

Link-State Versus Distance Vector

Link-state algorithms (also known as shortest path first algorithms) flood routing information to all nodes in the internetwork. Each router, however, sends only the portion of the routing table that describes the state of its own links. In link-state algorithms, each router builds a picture of the entire network in its routing tables. Distance vector algorithms (also known as Bellman-Ford algorithms) call for each router to send all or some portion of its routing table, but only to its neighbors. In essence, link-state algorithms send small updates everywhere, while distance vector algorithms send larger updates only to neighboring routers. Distance vector algorithms know only about their neighbors. Typically, link-state algorithms are used in conjunction with OSPF routing protocols.

Supported Internet Protocols for Routing

The Firewall Threat Defense device supports several Internet protocols for routing. Each protocol is briefly described in this section.

- Enhanced Interior Gateway Routing Protocol (EIGRP)

EIGRP is a Cisco proprietary protocol that provides compatibility and seamless interoperation with IGRP routers. An automatic-redistribution mechanism allows IGRP routes to be imported into Enhanced IGRP, and vice versa, so it is possible to add Enhanced IGRP gradually into an existing IGRP network.

- Open Shortest Path First (OSPF)

  OSPF is a routing protocol developed for Internet Protocol (IP) networks by the interior gateway protocol (IGP) working group of the Internet Engineering Task Force (IETF). OSPF uses a link-state algorithm to build and calculate the shortest path to all known destinations. Each router in an OSPF area includes an identical link-state database, which is a list of each of the router usable interfaces and reachable neighbors.

- Routing Information Protocol (RIP)

  RIP is a distance-vector protocol that uses hop count as its metric. RIP is widely used for routing traffic in the global Internet and is an interior gateway protocol (IGP), which means that it performs routing within a single autonomous system.

- Border Gateway Protocol (BGP)

  BGP is an interautonomous system routing protocol. BGP is used to exchange routing information for the Internet and is the protocol used between Internet service providers (ISP). Customers connect to ISPs, and ISPs use BGP to exchange customer and ISP routes. When BGP is used between autonomous systems (AS), the protocol is referred to as External BGP (EBGP). If a service provider is using BGP to exchange routes within an AS, then the protocol is referred to as Interior BGP (IBGP).

Routing Table

The Firewall Threat Defense uses separate routing tables for data traffic (through-the-device) and for management traffic (from-the-device). This section describes how the routing tables work. For information about the management routing table, see also Routing Table for Management Traffic, on page 1163.

How the Routing Table Is Populated

The Firewall Threat Defense routing table can be populated by statically defined routes, directly connected routes, and routes discovered by the dynamic routing protocols. Because the Firewall Threat Defense device can run multiple routing protocols in addition to having static and connected routes in the routing table, it is possible that the same route is discovered or entered in more than one manner. When two routes to the same destination are put into the routing table, the one that remains in the routing table is determined as follows:

- If the two routes have different network prefix lengths (network masks), then both routes are considered unique and are entered into the routing table. The packet forwarding logic then determines which of the two to use.

  For example, if the RIP and OSPF processes discovered the following routes:

  - RIP: 192.168.32.0/24
  - OSPF: 192.168.32.0/19

  Even though OSPF routes have the better administrative distance, both routes are installed in the routing table because each of these routes has a different prefix length (subnet mask). They are considered different destinations and the packet forwarding logic determines which route to use.

- If the Firewall Threat Defense device learns about multiple paths to the same destination from a single routing protocol, such as RIP, the route with the better metric (as determined by the routing protocol) is entered into the routing table.

  Metrics are values associated with specific routes, ranking them from most preferred to least preferred. The parameters used to determine the metrics differ for different routing protocols. The path with the lowest metric is selected as the optimal path and installed in the routing table. If there are multiple paths to the same destination with equal metrics, load balancing is done on these equal cost paths.

- If the Firewall Threat Defense device learns about a destination from more than one routing protocol, the administrative distances of the routes are compared, and the routes with lower administrative distance are entered into the routing table.

Administrative Distances for Routes

You can change the administrative distances for routes discovered by or redistributed into a routing protocol. If two routes from two different routing protocols have the same administrative distance, then the route with the lower *default* administrative distance is entered into the routing table. In the case of EIGRP and OSPF routes, if the EIGRP route and the OSPF route have the same administrative distance, then the EIGRP route is chosen by default.

Administrative distance is a route parameter that the Firewall Threat Defense device uses to select the best path when there are two or more different routes to the same destination from two different routing protocols. Because the routing protocols have metrics based on algorithms that are different from the other protocols, it is not always possible to determine the best path for two routes to the same destination that were generated by different routing protocols.

Each routing protocol is prioritized using an administrative distance value. The following table shows the default administrative distance values for the routing protocols supported by the Firewall Threat Defense device.

Table 70: Default Administrative Distance for Supported Routing Protocols

Route Source	Default Administrative Distance
Connected interface	0
VPN route	1
Static route	1
EIGRP Summary Route	5
External BGP	20
Internal EIGRP	90
OSPF	110
IS-IS	115
RIP	120
EIGRP external route	170
Internal and local BGP	200

Route Source	Default Administrative Distance
Unknown	255

The smaller the administrative distance value, the more preference is given to the protocol. For example, if the Firewall Threat Defense device receives a route to a certain network from both an OSPF routing process (default administrative distance - 110) and a RIP routing process (default administrative distance - 120), the Firewall Threat Defense device chooses the OSPF route because OSPF has a higher preference. In this case, the router adds the OSPF version of the route to the routing table.

A VPN advertised route (V-Route/RRI)) is equivalent to a static route with the default administrative distance 1. But it has a higher preference as with the network mask 255.255.255.255.

In this example, if the source of the OSPF-derived route was lost (for example, due to a power shutdown), the Firewall Threat Defense device would then use the RIP-derived route until the OSPF-derived route reappears.

The administrative distance is a local setting. For example, if you change the administrative distance of routes obtained through OSPF, that change would only affect the routing table for the Firewall Threat Defense device on which the command was entered. The administrative distance is not advertised in routing updates.

Administrative distance does not affect the routing process. The routing processes only advertise the routes that have been discovered by the routing process or redistributed into the routing process. For example, the RIP routing process advertises RIP routes, even if routes discovered by the OSPF routing process are used in the routing table.

Backup Dynamic and Floating Static Routes

A backup route is registered when the initial attempt to install the route in the routing table fails because another route was installed instead. If the route that was installed in the routing table fails, the routing table maintenance process calls each routing protocol process that has registered a backup route and requests them to reinstall the route in the routing table. If there are multiple protocols with registered backup routes for the failed route, the preferred route is chosen based on administrative distance.

Because of this process, you can create floating static routes that are installed in the routing table when the route discovered by a dynamic routing protocol fails. A floating static route is simply a static route configured with a greater administrative distance than the dynamic routing protocols running on the Firewall Threat Defense device. When the corresponding route discovered by a dynamic routing process fails, the static route is installed in the routing table.

How Forwarding Decisions Are Made

Forwarding decisions are made as follows:

- If the destination does not match an entry in the routing table, the packet is forwarded through the interface specified for the default route. If a default route has not been configured, the packet is discarded.
- If the destination matches a single entry in the routing table, the packet is forwarded through the interface associated with that route.
- If the destination matches more than one entry in the routing table, then the packet is forwarded out of the interface associated with the route that has the longer network prefix length.

For example, a packet destined for 192.168.32.1 arrives on an interface with the following routes in the routing table:

- 192.168.32.0/24 gateway 10.1.1.2
- 192.168.32.0/19 gateway 10.1.1.3

In this case, a packet destined to 192.168.32.1 is directed toward 10.1.1.2, because 192.168.32.1 falls within the 192.168.32.0/24 network. It also falls within the other route in the routing table, but 192.168.32.0/24 has the longest prefix within the routing table (24 bits verses 19 bits). Longer prefixes are always preferred over shorter ones when forwarding a packet.

Note Existing connections continue to use their established interfaces even if a new similar connection would result in different behavior due to a change in routes.

Dynamic Routing and High Availability

Dynamic routes are synchronized on the standby unit when the routing table changes on the active unit. This means that all additions, deletions, or changes on the active unit are immediately propagated to the standby unit. If the standby unit becomes active in an active/standby ready High Availability pair, it will already have an identical routing table as that of the former active unit because routes are synchronized as a part of the High Availability bulk synchronization and continuous replication processes.

Dynamic Routing in Clustering

The routing process only runs on the control node, and routes are learned through the control node and replicated to data nodes. If a routing packet arrives at a data node, it is redirected to the control node.

Figure 411: Dynamic Routing in Spanned EtherChannel Mode

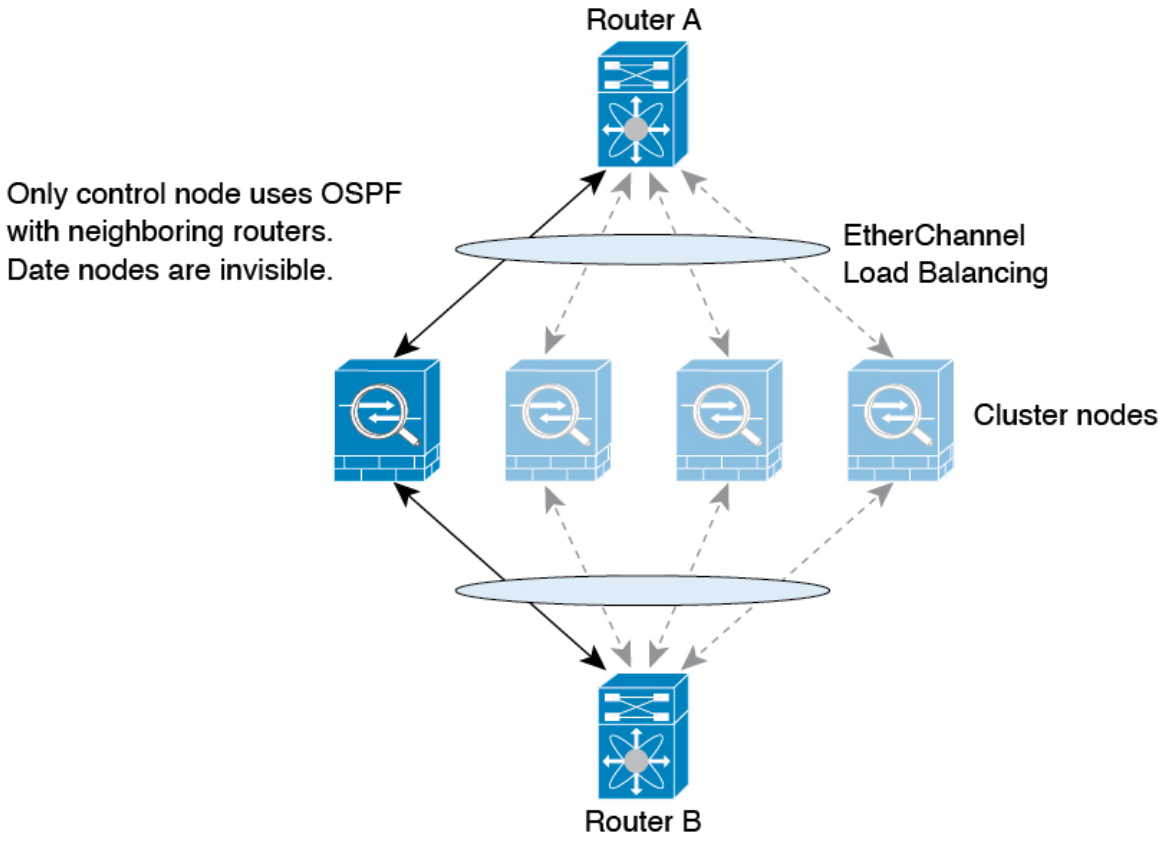

After the data node learn the routes from the control node, each node makes forwarding decisions independently.

The OSPF LSA database is not synchronized from the control node to data nodes. If there is a control node switchover, the neighboring router will detect a restart; the switchover is not transparent. The OSPF process picks an IP address as its router ID. Although not required, you can assign a static router ID to ensure a consistent router ID is used across the cluster. See the OSPF Non-Stop Forwarding feature to address the interruption.

Dynamic Routing in Individual Interface Mode

In Individual interface mode, each node runs the routing protocol as a standalone router, and routes are learned by each node independently.

Figure 412: Dynamic Routing in Individual Interface Mode

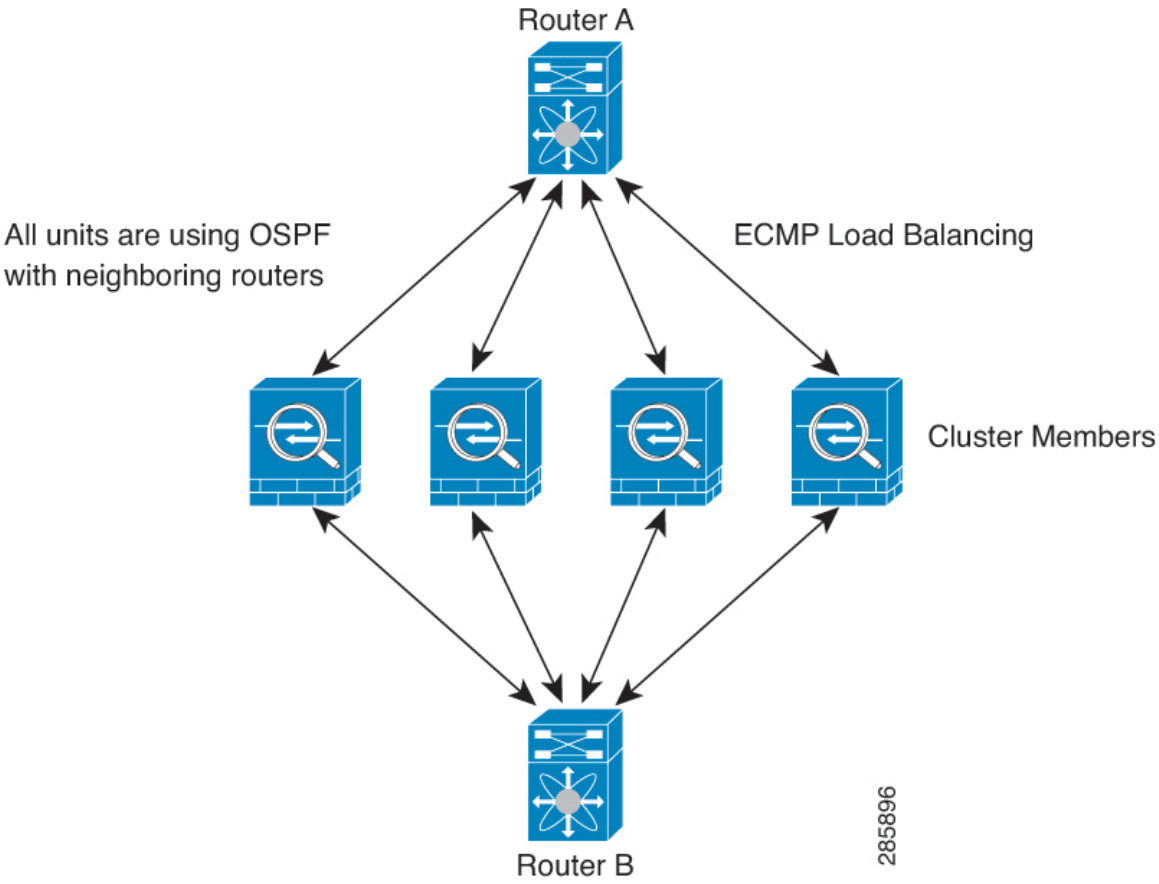

In the above diagram, Router A learns that there are 4 equal-cost paths to Router B, each through a node. ECMP is used to load balance traffic between the 4 paths. Each node picks a different router ID when talking to external routers.

You must configure a cluster pool for the router ID so that each node has a separate router ID.

EIGRP does not form neighbor relationships with cluster peers in individual interface mode.

Note If the cluster has multiple adjacencies to the same router for redundancy purposes, asymmetric routing can lead to unacceptable traffic loss. To avoid asymmetric routing, group all of these node interfaces into the same traffic zone. See Create an ECMP Zone, on page 1225.

Routing Table for Management Traffic

The Firewall Threat Defense device includes the following routing tables for from-the-device management traffic:

- Linux Management routing table—Special management traffic sourced from the Management interface such as Firewall Management Center communication,the licensing communication, and database updates always uses the Linux Management routing table.

- Data routing table—All from-the-device traffic (as well as all through traffic) uses the data routing table by default. All regular data interfaces are part of this routing table. Most services let you choose a specific interface, so only routes associated with that interface are used.
- Management-only routing table—The Management interface and all data interfaces that you set to management-only are part of this routing table. To send from-the-device traffic from any of these interfaces, you must choose a specific management-only interface when you configure the service. An exception is for DNS lookups and ICMP (ping and traceroute): in these cases, the Firewall Threat Defense will use data and then fall back to management automatically if a route is not found. You can add static routes for management-only interfaces, but not for the special Management interface. The Firewall Threat Defense device automatically adds a default route for Management that forwards traffic to Linux, where a separate route lookup occurs in the Linux routing table. You can add static routes to the Linux routing table that can be used by Management using the Firewall Threat Defense CLI **configure network static-routes** command.

Note The *default* Linux route is set with the **configure network ipv4** or **configure network ipv6** command.

Note For devices that have not yet merged the Management and legacy Diagnostic interfaces, see refer to pre-7.3 versions of this guide.

Equal-Cost Multi-Path (ECMP) Routing

The Firewall Threat Defense device supports Equal-Cost Multi-Path (ECMP) routing.

You can have up to 8 equal cost static or dynamic routes per interface. For example, you can configure multiple default routes on the outside interface that specify different gateways.

```
route for 0.0.0.0 0.0.0.0 through outside to 10.1.1.2
route for 0.0.0.0 0.0.0.0 through outside to 10.1.1.3
route for 0.0.0.0 0.0.0.0 through outside to 10.1.1.4
```

In this case, traffic is load-balanced on the outside interface between 10.1.1.2, 10.1.1.3, and 10.1.1.4. Traffic is distributed among the specified gateways based on an algorithm that hashes the source and destination IP addresses, incoming interface, protocol, source and destination ports.

ECMP Across Multiple Interfaces Using Traffic Zones

If you configure traffic zones to contain a group of interfaces, you can have up to 8 equal cost static or dynamic routes across up to 8 interfaces within each zone. For example, you can configure multiple default routes across three interfaces in the zone:

```
route for 0.0.0.0 0.0.0.0 through outside1 to 10.1.1.2
route for 0.0.0.0 0.0.0.0 through outside2 to 10.2.1.2
route for 0.0.0.0 0.0.0.0 through outside3 to 10.3.1.2
```

Similarly, your dynamic routing protocol can automatically configure equal cost routes. The Firewall Threat Defense device load-balances traffic across the interfaces with a more robust load balancing mechanism.

When a route is lost, the device seamlessly moves the flow to a different route.

About Route Maps

Route maps are used when redistributing routes into an OSPF, RIP, EIGRP or BGP routing process. They are also used when generating a default route into an OSPF routing process. A route map defines which of the routes from the specified routing protocol are allowed to be redistributed into the target routing process.

Route maps have many features in common with widely known ACLs. These are some of the traits common to both:

- They are an ordered sequence of individual statements, and each has a permit or deny result. Evaluation of an ACL or a route map consists of a list scan, in a predetermined order, and an evaluation of the criteria of each statement that matches. A list scan is aborted once the first statement match is found and an action associated with the statement match is performed.
- They are generic mechanisms. Criteria matches and match interpretation are dictated by the way that they are applied and the feature that uses them. The same route map applied to different features might be interpreted differently.

These are some of the differences between route maps and ACLs:

- Route maps are more flexible than ACLs and can verify routes based on criteria which ACLs can not verify. For example, a route map can verify if the type of route is internal.
- Each ACL ends with an implicit deny statement, by design convention. If the end of a route map is reached during matching attempts, the result depends on the specific application of the route map. Route maps that are applied to *redistribution* behave the same way as ACLs: if the route does not match any clause in a route map then the route redistribution is denied, as if the route map contained a deny statement at the end.

Permit and Deny Clauses

Route maps can have permit and deny clauses. The deny clause rejects route matches from redistribution. You can use an ACL as the matching criterion in the route map. Because ACLs also have permit and deny clauses, the following rules apply when a packet matches the ACL:

- ACL permit + route map permit: routes are redistributed.
- ACL permit + route map deny: routes are not redistributed.
- ACL deny + route map permit or deny: the route map clause is not matched, and the next route-map clause is evaluated.

Match and Set Clause Values

Each route map clause has two types of values:

- A match value selects routes to which this clause should be applied.
- A set value modifies information that will be redistributed into the target protocol.

For each route that is being redistributed, the router first evaluates the match criteria of a clause in the route map. If the match criteria succeeds, then the route is redistributed or rejected as dictated by the permit or deny

clause, and some of its attributes might be modified by the values set from the set commands. If the match criteria fail, then this clause is not applicable to the route, and the software proceeds to evaluate the route against the next clause in the route map. Scanning of the route map continues until a clause is found that matches the route or until the end of the route map is reached.

A match or set value in each clause can be missed or repeated several times, if one of these conditions exists:

- If several match entries are present in a clause, all must succeed for a given route in order for that route to match the clause (in other words, the logical AND algorithm is applied for multiple match commands).
- If a match entry refers to several objects in one entry, either of them should match (the logical OR algorithm is applied).
- If a match entry is not present, all routes match the clause.
- If a set entry is not present in a route map permit clause, then the route is redistributed without modification of its current attributes.

Note Do not configure a set entry in a route map deny clause because the deny clause prohibits route redistribution—there is no information to modify.

A route map clause without a match or set entry does perform an action. An empty permit clause allows a redistribution of the remaining routes without modification. An empty deny clause does not allow a redistribution of other routes (this is the default action if a route map is completely scanned, but no explicit match is found).

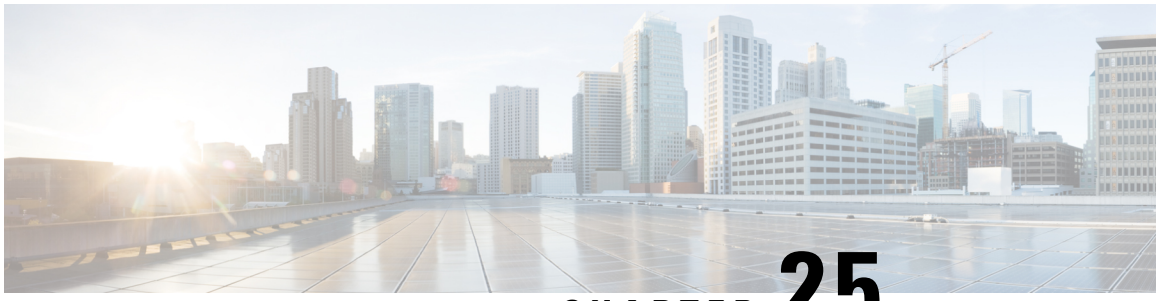

CHAPTER 25

Virtual Routers

This chapter describes underlying concepts about virtual routers and on how virtual routing behaves within the Secure Firewall Threat Defense.

- About Virtual Routers and Virtual Routing and Forwarding (VRF), on page 1167
- Maximum Number of Virtual Routers By Device Model, on page 1174
- Requirements and Prerequisites for Virtual Routers, on page 1175
- Guidelines and Limitations for Virtual Routers, on page 1175
- Modifications to the Firewall Management Center Web Interface - Routing Page, on page 1177
- Manage Virtual Routers, on page 1178
- Create a Virtual Router, on page 1178
- Monitoring Virtual Routers, on page 1182
- Configuration Examples for Virtual Routers, on page 1182
- History for Virtual Routers, on page 1221

About Virtual Routers and Virtual Routing and Forwarding (VRF)

You can create multiple virtual routers to maintain separate routing tables for groups of interfaces. Because each virtual router has its own routing table, you can provide clean separation in the traffic flowing through the device.

Thus, you can provide support to two or more distinct customers over a common set of networking equipment. You can also use virtual routers to provide more separation for elements of your own network, for example, by isolating a development network from your general purpose corporate network.

Virtual routers implement the “light” version of Virtual Routing and Forwarding, or VRF-Lite, which does not support Multiprotocol Extensions for BGP (MBGP).

When you create a virtual router, you assign interfaces to the router. You can assign a given interface to one, and only one, virtual router. You would then define static routes, and configure routing protocols such as OSPF or BGP, for each virtual router. You would also configure separate routing processes over your entire network, so that routing tables on all participating devices are using the same per-virtual-router routing process and tables. Using virtual routers, you create logically-separated networks over the same physical network to ensure the privacy of the traffic that runs through each virtual router.

Because the routing tables are separate, you can use the same, or overlapping, address spaces across the virtual routers. For example, you could use the 192.168.1.0/24 address space for two separate virtual routers, supported by two separate physical interfaces.

Note that there are separate management and data routing tables per virtual router. For example, if you assign a management-only interface to a virtual router, then the routing table for that interface is separate from the data interfaces assigned to the virtual router.

About Virtual Routers and Dynamic VTI

Virtual Routers and Dynamic VTI

You can create virtual routers, associate dynamic VTIs with these virtual routers, and extend the capabilities of dynamic VTIs in your network. You can associate dynamic VTIs either with global or user-defined virtual routers. You can assign a dynamic VTI to only one virtual router.

A virtual router associated with:

- A dynamic VTI is called an Indoor VRF (IVRF).
- A tunnel source interface is known as Front Door VRF (FVRF).

A dynamic VTI and its corresponding protected network interface must be part of the same virtual router. You must map the borrow IP interface and the dynamic VTI to the same virtual router. A tunnel source interface can be part of multiple virtual routers.

To configure virtual routers using dynamic VTI for a route-based site-to-site VPN, see How to Configure a Virtual Router with Dynamic VTI, on page 1168.

For more information about a configuration example, see How to Secure Traffic from Networks with Multiple Virtual Routers over a Site-to-Site VPN with Dynamic VTI, on page 1201

How to Configure a Virtual Router with Dynamic VTI

To configure a virtual router with dynamic VTI for a route-based site-to-site VPN in the management center:

Step	Do This	More Information
1	Create a route-based site-to-site VPN with a dynamic VTI interface on the hub and static VTI(s) on the spoke(s).	Create a Route-based Site-to-Site VPN, on page 1538
2	Create a virtual router.	Create a Virtual Router, on page 1178
3	Assign the interfaces to the virtual router.	Configure a Virtual Router, on page 1179
4	Configure routing policies for the hub and spoke.	Configure Endpoints for a Hub and Spoke Topology, on page 1542
5	Configure access control policies for the hub and spoke.	Configure Endpoints for a Hub and Spoke Topology, on page 1542

Applications of Virtual Routers

You can use virtual routers to isolate network on shared resources and/or to isolate networks with common security policy. Thus, virtual routers help you to achieve:

- Traffic separation for customers through dedicated routing tables for each customer or for different departments.
- Common security policy management for different departments or networks.
- Shared internet access for different departments or network.

Global and User-Defined Virtual Routers

Global Virtual Routers

For a device with virtual routing capability, system creates a global virtual router by default. The system assigns all interfaces in your network to the global virtual router. A routed interface can belong to either a user-defined virtual router or a global virtual router. When you upgrade Firewall Threat Defense to a version which has virtual router capability, all its existing routing configuration becomes part of the global virtual router.

User-Defined Virtual Routers

A user-defined virtual router is the one defined by you. You can create more than one virtual router on a device. However, anytime, an interface can be assigned to only one user-defined virtual router. While some of the device features are supported on user-defined virtual routers, few of the features are supported only on the global virtual routers. User-defined virtual routers support route-based site-to-site VPN (static VTI) (static and dynamic VTI).

Supported Features and Monitoring Policies

You can configure the following features on the global virtual router only:

- OSPFv3
- RIP
- EIGRP
- IS-IS
- Multicast Routing
- Policy Based Routing (PBR)

ISIS and PBR are supported through Flex Config in Firewall Management Center (see Predefined FlexConfig Objects, on page 2678). Configure only global virtual router interfaces for these features.

DHCP server auto-configuration uses WINS/DNS server that is learned from an interface. This interface can only be a global virtual router interface.

You can configure the following features separately for each user-defined virtual router:

- Static routes and their SLA monitors
- OSPFv2
- BGPv4/v6
- Integrated Routing and Bridging (IRB)

- SNMP

Following features are used by the system when querying or communicating with the remote system (from-the-box traffic). These features use interfaces in the global virtual router only. That means, if you configure an interface for the feature, it must belong to the global virtual router. As a rule, anytime, if the system must look up a route to reach an external server for its own management purposes, it does the route lookup in the global virtual router.

- DNS server when used to resolve fully qualified names used in access control rules, or for resolving names for the **ping** command. If you specify **any** as the interface for a DNS server, the system considers interfaces in the global virtual router only.
- AAA server or identity realm when used with VPN. You can configure VPN on interfaces in the global virtual router only. Thus, the external AAA servers that are used for VPN, such as Active Directory, must be reachable through an interface in the global virtual router.

Configuring Policies to be Virtual-Router-Aware

When you create a virtual router, the routing table for that virtual router is automatically separated from the global virtual router or any other virtual router. However, security policies are not automatically virtual-router-aware.

For example, if you write an access control rule that applies to “any” source or destination security zone, then the rule will apply to all interfaces across all virtual routers. This might in fact be exactly what you want. For example, all of your customers might want to block access to the same list of objectionable URL categories.

But, if you need to apply a policy to one of the virtual routers but not others, you need to create security zones that contain interfaces from that single virtual router only. Then, use the virtual-router-constrained security zones in the source and destination criteria of the security policy.

By using security zones whose memberships are constrained to the interfaces assigned to a single virtual router, you can write virtual-router-aware rules in the following policies:

- Access control policy.
- Intrusion and file policies.
- SSL decryption policy.
- Identity policy and user-to-IP address mappings. If you use overlapping address spaces in virtual routers, ensure that you create separate realms for each virtual router and apply them correctly in the identity policy rules.

If you use overlapping address spaces in your virtual routers, you should use security zones to ensure that the right policies get applied. For example, if you use the 192.168.1.0/24 address space in two separate virtual routers, an access control rule that simply specifies the 192.168.1.0/24 network will apply to traffic in both virtual routers. If that is not the desired outcome, you can limit the application of the rule by also specifying the source/destination security zones for just one of the virtual routers.

Interconnecting Virtual Routers

Static and Dynamic Route Leaking

You can configure the device to route traffic between virtual routers. This process of route leaking can be done manually by setting up static routes or dynamically through BGP settings.

Static Route Leaking

You can configure static routes to route traffic between virtual routers.

For example, if you have the outside interface in the global virtual router, you can set up static default routes in each of the other virtual routers to send traffic to the outside interface. Then, any traffic that cannot be routed within a given virtual router gets sent to the global router for subsequent routing.

Static routes between virtual routers are known as route leaks, because you are leaking traffic to a different virtual router. When you are leaking routes, say, VR1 routes to VR2, you can initiate connections from VR2 to VR1 only. For traffic to flow from VR1 to VR2, you must configure the reverse route. When you create a static route to an interface in another virtual router, you do not need to specify a gateway address. Simply select the destination interface.

For inter-virtual-router routes, the system does destination interface look-up in the source virtual router. Then, it looks up the MAC address of the next hop in the destination virtual router. Thus, the destination virtual router must have either a dynamic (learned) or static route for the selected interface for the destination address.

Configuring NAT rules that use source and destination interfaces in different virtual routers can also allow traffic to route between virtual routers. If you do not select the option for NAT to do a route lookup, the rule will simply send traffic out the destination interface with a NATed address whenever destination translation happens. However, the destination virtual router should have a route for the translated destination IP address so that next-hop lookup can succeed.

Though NAT rule leaks traffic from one virtual router to another, to ensure correct routing, we recommend that you configure a static route leak between these virtual routers for the translated traffic. Without the route leak, sometimes the rule may not match the traffic you expect it to match, and the translation may not be applied.

Virtual routing does not support a cascading or chain of route leaks. For example, assume that your Firewall Threat Defense has VR1, VR2, and VR3 virtual routers; VR3 is directly connected to a network - 10.1.1.0/24. Now, assume you configure a route leak in VR1 for network 10.1.1.0/24 through interface in VR2 and in VR2 define a route leak for 10.1.1.0/24 through VR3. This chain of route leaks will not allow traffic to hop from VR1 to VR2 and then exit from VR3. In case of route leaks, the route lookups first determine egress interface from input Virtual Router’s routing table and then looks at the output of Virtual Router’s routing table for next hop lookup. From both the lookups, egress interface should match. In our example, the egress interfaces will not be the same and hence the traffic does not pass through.

Use static inter VRF route with caution when the destination network is not a direct-connected subnet of the upstream (outgoing) VR. For example, assume two VRs - VR1 and VR2. While VR1 handles the outgoing traffic which gets the default route from its external peer through BGP or any dynamic routing protocol, and VR2 handles the incoming traffic which is configured with static inter VRF default route with VR1 as the next-hop. When VR1 loses the default route from its peer, VR2 will not able to detect that its upstream (outgoing) VR lost the default route and the traffic is still sent toward VR1 which will eventually get dropped without notifications. In this scenario, we recommend that you configure VR2 with dynamic route leak through BGP.

Dynamic Route Leaking Using BGP

You can implement an inter-virtual-router route leak by exporting routes from a source virtual router (say VR1) to the source BGP table using route target extended community and then importing the same route target extended community from the source BGP table into the destination BGP table, which in turn is used by the destination virtual router (say, VR2). You can use the route maps for filtering the routes. The routes of global virtual router can also be leaked to user-defined virtual routers and vice versa. The BGP inter-virtual-router route leaking supports both ipv4 and ipv6 prefixes.

For details on configuring BGP route leaking, see Configure BGP Route Import/Export Settings, on page 1293.

BGP Route Leaking Guidelines

- Ensure that all the routes required for recursion are imported and present in the routing table of the ingress virtual router.
- ECMP is supported per virtual router. Hence, do not configure an ECMP across different virtual routers. The overlapping prefixes imported from different virtual routers cannot form an ECMP. That is, when you attempt to import routes with overlapping addresses from two different virtual routers to other virtual routers (a global virtual router or an user-defined virtual router), only one route (as per BGP best path algorithm, the first route that was advertised) is imported to the respective virtual routing table. For example, if a network 10.10.0.0/24, connected to VR1 is advertised through BGP to a global virtual router first, and later another network with the same address 10.10.0.0/24, connected to VR2 is also advertised through BGP to global virtual router, only the VR1 network route is imported to the global virtual routing table.
- OSPFv3 is not supported on user-defined virtual routers. Hence, do not configure BGPv6 to leak OSPFv3 user-defined virtual routers to global virtual router. However, you can configure BGPv6 to leak OSPFv3 global virtual router routes to user-defined virtual router through redistribution.
- It is recommended to keep VTI interface, protected internal interfaces (loopback interface if supported for VTI) to be part of same virtual router to prevent the need for route leak.

Overlapping IP Addresses

Virtual router creates multiple instances of routing tables that are independent, thereby, the same or overlapping IP addresses can be used without conflicts. Firewall Threat Defense allows the same network to be part of two or more virtual routers. This involves multiple policies to be applied at the interface or at the virtual router level.

Other than few exceptions, the routing functions and most of the NGFW and IPS capability does not get impacted by the overlapping IP addresses. The following section describes the features that have limitations with overlapping IP addresses and the suggestions or recommendations to overcome them.

Limitations with Overlapping IP Addresses

When using an overlapping IP address in multiple virtual routers, to ensure proper application of the policy, you have to modify policies or rules for some of the features. Such features require you to use more specific interface either by splitting existing security zone or using new interface group as the need be.

Following features need modification for its proper functioning with an overlapping IP address:

- Network Map—Modify the network discovery policy to exclude some overlapping IP segments to ensure that there is no overlapping IP address being mapped.

- Identity Policy—The identity feed source cannot differentiate among virtual routers; to overcome this limitation, map overlapping address spaces or virtual routers in different realms.

For the following features, you need to apply rules on specific interfaces to ensure that different policies are applied on overlapping IP segments:

- Access Policy
- Prefilter Policy
- QoS/Rate Limit
- SSL Policy

Unsupported Features with Overlapping IP Addresses

- ISE SGT-based Rule in AC Policy—The static security group tag (SGT) to IP address mappings downloaded from Cisco Identity Services Engine (ISE) are not virtual-router-aware. Set up separate ISE systems per virtual router if you need to create different SGT mappings per virtual router. This is not necessary if you intend to map the same IP addresses to the same SGT number in each virtual router.
- Overlapping DHCP server pools are not supported across virtual routers.
- Events and Analytics—Many of the Firewall Management Center analytics are dependent on network map and identity mappings which cannot differentiate if the same IP address belongs to two different end hosts. Hence, these analytics are not accurate when there are overlapping IP segments existing in same device but different virtual routers.

Configuring SNMP on User-Defined Virtual Routers

In addition to supporting SNMP on the management interface and Global virtual router data interfaces, Secure Firewall Threat Defense now allows you to configure SNMP host on user-defined virtual routers.

Configuring an SNMP host on user-defined virtual routers includes the following process:

1. Configure device interfaces.
2. Configure virtual router on an interface.
3. Configure SNMP hosts on a virtual router interface.

Note SNMP is not virtual router-aware. Hence, while configuring SNMP server on the user-defined virtual router, ensure that the network address is not an overlapping IP address.

4. Deploy configurations to Secure Firewall Threat Defense. On successful deployment, the-SNMP polling and traps are sent to the Network Management Station through the virtual router interface.

Maximum Number of Virtual Routers By Device Model

The maximum number of virtual routers you can create depends on the device model. The following table provides the maximum limits. You can double-check on your system by entering the **show vrf counters** command, which shows the maximum number of user-defined virtual routers for that platform not including the global virtual router. The numbers in the table below include user and global routers. For the Firepower 4100/9300, these numbers apply to native mode.

For platforms that support multi-instance capability, such as the Firepower 4100/9300, determine the maximum number of virtual routers per container instance by dividing the maximum virtual routers by the number of cores on the device, and then multiplying by the number of cores assigned to the instance, rounding down to the nearest whole number. For example, if the platform supports a maximum of 100 virtual routers, and it has 70 cores, then each core would support a maximum of 1.43 virtual routers (rounded). Thus, an instance assigned 6 cores would support 8.58 virtual routers, which rounds down to 8, and an instance assigned 10 cores would support 14.3 virtual routers (rounding down, 14).

Device Model	Maximum Virtual Routers
Firepower 1010	5
Firepower 1120	5
Firepower 1140	10
Firepower 1150	10
Secure Firewall 1210CE	5
Secure Firewall 1210CP	5
Secure Firewall 1220CX	10
Secure Firewall 1230	10
Secure Firewall 1240	10
Secure Firewall 1250	15
Secure Firewall 3105	10
Secure Firewall 3110	15
Secure Firewall 3120	25
Secure Firewall 3130	50
Secure Firewall 3140	100
Firepower 4112	60
Firepower 4115	80
Firepower 4125	100

Device Model	Maximum Virtual Routers
Firepower 4145	100
Secure Firewall 4215	100
Secure Firewall 4225	100
Secure Firewall 4245	100
Firepower 9300 appliance, all models	100
Firewall Threat Defense Virtual, all platforms	30
ISA 3000	10

Related Topics

Requirements and Prerequisites for Container Instances, on page 323

Requirements and Prerequisites for Virtual Routers

Model Support

Threat Defense

Supported Domains

Any

User Roles

Admin

Network Admin

Security Approver

Guidelines and Limitations for Virtual Routers

Firewall Mode Guidelines

Virtual routers are supported on routed firewall mode only.

Interface Guidelines

- You can assign an interface to only one virtual router.
- A virtual router can have any number of interfaces that are assigned to it.

- You can assign only routed interfaces with logical names and VTIs to a user-defined virtual router.
- If you want to change a virtual router interface to a non-routed mode, remove the interface from the virtual router, and then change its mode.
- You can assign an interface to a virtual router, either from a global virtual router or from another user-defined virtual router.
- The following interfaces cannot be assigned to an user-defined virtual router:
  - Members of EtherChannel.
  - Members of Redundant interfaces.
  - Members of BVI.
- VTI is a route-based VPN. So, when the tunnel is established, the traffic that uses VTI for encryption must be controlled through routing. Static routing, as well as dynamic routing with BGP, OSPFv2/v3, or EIGRP is supported.
- You cannot use interfaces that belong to user-defined virtual routers in policy-based site-to-site or remote access VPNs.
- A dynamic VTI and its corresponding protected network interface must be part of the same virtual router.
- You must map the borrow IP interface and the dynamic VTI to the same virtual router.
- User-defined virtual routers support only BGPv4/v6 and OSPFv2 routing protocols.
- A tunnel source interface can be in a different user-defined virtual router than that associated with the dynamic VTI.
- If a route using the interface that is being moved or its virtual router is deleted, exist in source or destination virtual router table, remove the routes before the interface movement or virtual router deletion.
- As separate routing tables are maintained for each virtual router, when an interface is moved from one virtual router to another virtual router, be it global or user-defined, the system removes the IP address configured on the interface temporarily. All existing connections on the interface are terminated. Thus, moving interfaces between virtual routers have drastic effect on the network traffic. Hence take precautionary measures before you move interfaces.

Global Virtual Router Guidelines

- The interfaces which are named and not part of other virtual routers, are part of the global virtual router.
- You cannot remove routed interfaces from global virtual router.
- You cannot modify global virtual router.
- Generally, after configuring interfaces, if you un-register and register back to same or another Firewall Management Center, interface configuration is imported back from device. With virtual router support, there is a restriction—the IP address for only global virtual router interfaces is retained.

Clustering Guidelines

- When the control unit link fails due to failure of its interfaces, the unit removes all leaked routes of its interfaces from the global routing table and propagates the inactive connected and static routes to other units of the cluster. This results in removal of those leaked routes from the routing table in other units. These removal takes place prior to another unit becoming a new control unit, which takes approximately 500 ms. When another unit becomes the new control unit, these routes are learned and added back to the routing tables through BGP convergence. Thus, till the convergence time, approximately one minute, the leaked routes are not available for the routing events to take place.
- When a control role change occurs in a cluster, the leaked routes learnt through BGP is updated with the best ECMP path. However, the non-best ECMP path is removed from the cluster routing table only after the BGP reconvergence timer elapses, that is 210 seconds. Thus, till the BGP reconvergence timer elapse, the old, non-best ECMP path persist as preferred route for routing events.

Additional Guidelines

- While configuring BGP for virtual routers, you can redistribute the routes belonging to different protocols within the same virtual routers. For example, OSPF VR2 routes cannot be imported into BGP VR1. You can only redistribute OSPF VR2 into BGP VR2, and then configure a route leak between BGP VR2 and BGP VR1.
- You cannot use IPv6 ACL for filtering the routes in the route map. Only prefix list is supported.
- Security Intelligence Policy—The Security Intelligence policy is not virtual-router-aware. If you add an IP address, URL, or DNS name to the block list, it is blocked for all virtual routers. This limitation is applicable on the interface having security zones.
- NAT Rules—Do not mix interfaces in NAT rules. In virtual routing, if the specified source and destination interface objects (interface groups or security zones) have interfaces that belong to different virtual routers, the NAT rule diverts traffic from one virtual router through another virtual router. The NAT does the route lookup in the virtual router table for the inbound interface only. If necessary, define static routes in the source virtual router for the destination interface. If you leave the interface as **any**, the rule applies to all interfaces regardless of virtual router membership.
- DHCP Relay—Interconnecting virtual routers for DHCP relay is not supported. For example, if DHCP relay client is enabled on VR1 interface and DHCP relay server is enabled on VR2 interface, the DHCP requests will not be forwarded outside of VR2 interface.
- Recreating a deleted virtual router—When you recreate a virtual router that was deleted less than 10 seconds earlier, an error message appears stating that deletion of the virtual router is in progress. If you want to recreate a deleted virtual router successively, use a different name for the new virtual router.

Modifications to the Firewall Management Center Web Interface - Routing Page

Devices earlier to the Firewall Threat Defense 6.6 and few device models are not supported with virtual routing capability. The Firewall Management Center web interface displays the same Routing page of the Firewall Management Center 6.5 or earlier version for such non supported devices. To know the supported devices and platform for virtual routing, see Supported Device Models.

You can configure virtual routers in the routing page of a supported device:

1. Navigate to **Devices** > **Device Management** and edit the virtual-router-aware device.
2. Click **Routing** to enter the virtual routers page.

For devices using virtual routing, the left pane of the Routing page displays the following:

- **Manage Virtual Routers**—allows you to create and manage virtual routers.
- List of virtual routing protocols—lists routing protocols that you can configure for the virtual routers.
- **General Settings**—allows you to configure BGP general settings that are applicable for all the virtual routers. Select the **Enable BGP** check box in order to define other BGP settings. To configure other BGP settings for a virtual router, navigate to **BGP** in the virtual routing protocols .

Manage Virtual Routers

When you click **Manage Virtual Routers** on the Virtual Routers pane, the Manage Virtual Routers page appears. This page displays the existing virtual routers on the device and the associated interfaces. In this page, you can **Add Virtual Router** (+) to the device. You can also **Edit** (✎) and **Delete** (🗑) user-defined virtual routers. You cannot edit or remove a global virtual router. You can only **View** (🔍) the details of a global virtual router.

Create a Virtual Router

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **Manage Virtual Routers**.

Step 4 Click **Add Virtual Router** (+).

Step 5 In the Add Virtual Router box, enter a name and description for the virtual router.

Note
If you are creating a virtual router that was deleted less than 10 seconds earlier, an error message appears stating that deletion of the virtual router is in progress. If you want to create a deleted virtual router successively, use a different name for the new virtual router.

Step 6 Click **Ok**.

The Routing page appears, displaying the newly created virtual router page.

What to do next

- Configure a Virtual Router.

Configure a Virtual Router

You can assign interfaces to a user-defined virtual router and configure the routing policies for the device. Though you cannot manually add or remove interfaces for a global virtual router, you can configure the routing policies for the device interfaces.

Before you begin

- To configure routing policies for a user-defined virtual router, add a router. See Create a Virtual Router, on page 1178.
- All routing configuration settings of a non-virtual routing capable device are also available for a global virtual router. For information on the settings, see Routing Settings.
- Only limited routing protocols are supported for a user-defined virtual router.

Procedure

Step 1 From the **Devices** > **Device Management** page, edit the virtual-router supported device. Navigate to **Routing**. For information on the modifications to the routing page, see Modifications to the Firewall Management Center Web Interface - Routing Page, on page 1177.

Step 2 From the drop-down list, select the desired virtual router.

Step 3 In the **Virtual Router Properties** page, you can modify the description.

Step 4 To add interfaces, select the interface under the **Available Interfaces** box, and then click **Add**.

Remember the following:

- Only interfaces with a logical name are listed under the **Available Interfaces** box. You can edit the interface and provide a logical name in **Interfaces**. Remember to save the changes for the settings to take effect.
- Only interfaces of global virtual routers are available for assigning; the **Available Interfaces** box lists only interfaces that are not assigned to any other user-defined virtual routers. You can assign physical interfaces, subinterfaces, redundant interfaces, bridge groups, VTIs, and EtherChannels to a virtual router, but not their member interfaces. Because the member interfaces cannot be named, they cannot be used in virtual routing.

  You can assign the diagnostic interface to the global virtual router only.

Step 5 To save the settings, click **Save**.

Step 6 To configure the routing policy for the virtual router, click the respective names to open the corresponding settings page:

- **OSPF**—Only OSPFv2 is supported on the user-defined virtual router. All other settings for OSPFv2 are as applicable as for a non-virtual-router-aware interface, except that **Interface** allows you to select only

the interfaces of the virtual router that you are configuring. You can define the OSPFv3 and OSPFv2 routing policies for a global virtual router. For information on the OSPF settings, see OSPF, on page 1239.

- **RIP**—You can configure RIP routing policies only for a global virtual router. For information on RIP settings, see RIP, on page 1297.
- **BGP**—This page displays the BGP general settings that you have configured in **Settings**:
  - You cannot modify any of those general settings on this page, except for the router ID settings. You can override the router ID settings that were defined in the **Settings** page by editing them on this page.
  - To configure other BGP IPv4 or IPv6 settings, you must enable the BGP option in **BGP** page under **General Settings**.
  - BGP configuration for both IPv4 and IPv6 address family are supported for global router and the user-defined virtual router.

  For information on configuring BGP settings, see BGP, on page 1277.
- **Static Route**—Use this setting to define where to send traffic for a specific destination network. You can also use this setting to create an inter-virtual-router static route. You can create a leak of connected or static route by using the interfaces of user-defined or global virtual routers. **FMC prefixes** to an interface to indicate that it is belonging to another virtual router and can be used for a route leak. For the route leak to be successful, do not specify next hop gateway.

  The Static Route table displays the virtual router whose interface is used for a route leak in the **Leaked from Virtual Router** column. If it is not a route leak, the column displays N/A.

  Irrespective of which virtual router the static route belongs, a Null0 interface is listed along with the interfaces of the same virtual router to which the static route belongs.

  For information on static route settings, see Static and Default Routes, on page 1151.
- **Multicast**—You can configure multicast routing policies only for a global virtual router. For information on multicast settings, see Multicast, on page 1303.

Step 7 To save the settings, click **Save**.

What to do next

- Modify a virtual router.
- Remove a virtual router.

Modify a Virtual Router

You can modify the description and other routing policies of a virtual router.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **Manage Virtual Routers**.

All virtual routers along with the assigned interfaces are displayed in the **Virtual Routers** page.

Step 4 To modify a virtual router, click **Edit** (✎) against the desired virtual router.

Note
You cannot modify the general settings of the global virtual router. Hence, edit is not available for the global router; instead a **View** (🔍) is provided to view the settings.

Step 5 To save the changes, click **Save**.

What to do next

- Remove a virtual router.

Remove Virtual Routers

Before you begin

- You cannot delete the Global virtual router. Hence, the delete option is not available for the Global virtual router.
- You can remove multiple virtual routers at a time.
- All the routing policies of the deleted virtual router are also deleted.
- All the interfaces of the deleted virtual router move to the global virtual router.
- If there are any restrictions on the movement of interfaces, such as overlapping IPs, route conflicts, and so on, you can remove the router only after resolving the conflicts.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **Manage Virtual Routers**.

All virtual routers along with the mapped interfaces are displayed in the **Virtual Routers** page.

Step 4 To remove a virtual router, click **Delete** (🗑) against the desired virtual router.

Step 5 To remove multiple routers, holding the CTRL key, click the virtual routers that you want to delete. Right-click, and then click **Delete**.

Step 6 To save the changes, click **Save**.

Monitoring Virtual Routers

To monitor and troubleshoot virtual routers, log into the device CLI and use the following commands:

- **show vrf**: Displays the details of the virtual routers and their associated interfaces.
- **show route vrf** <*vrf_name*>: Displays the routing details of a virtual router.
- **show run router bgp all**: Displays the BGP routing details of all virtual router.
- **show run router bgp vrf** <*vrf_name*>: Displays the BGP routing details of a virtual router.
- **show crypto ipsec sa/show crypto ikev2 sa**: Displays the details of the tunnel and the associated virtual router.
- You can monitor the tunnels in the Site-to site monitoring dashboard (**Overview > Site to Site VPN**).

  In the **Tunnel Status** widget, hover over a topology, click View and click **Packet Tracer** to view and troubleshoot the threat defense VPN tunnels.

Configuration Examples for Virtual Routers

How to Route to a Distant Server through Virtual Routers

In virtual routing, you can create multiple virtual routers to maintain separate routing tables for groups of interfaces, thereby achieve network separation. In some scenarios, you may need to access a server that is reachable only through a separate virtual router. This example provides the procedure that interconnects virtual routers to reach to a host that is multiple hops away.

Consider an example, where a member of the sales department of a garment company wants to look up at the stock maintained by the warehousing department of its factory unit. In a virtual routing environment, you need to leak route between virtual routers where destination (warehousing department) is multiple hops away from sales department. This route leaking is done by adding multihop route leak, where, you configure a static route in Sales virtual router(source) to an interface in Warehouse virtual router (destination). As the destination network is multi-hop away, you also need to configure the Warehouse virtual router with the route to the destination network, namely 10.50.0.0/24.

Figure 413: Interconnecting Two Virtual Routers - An Example

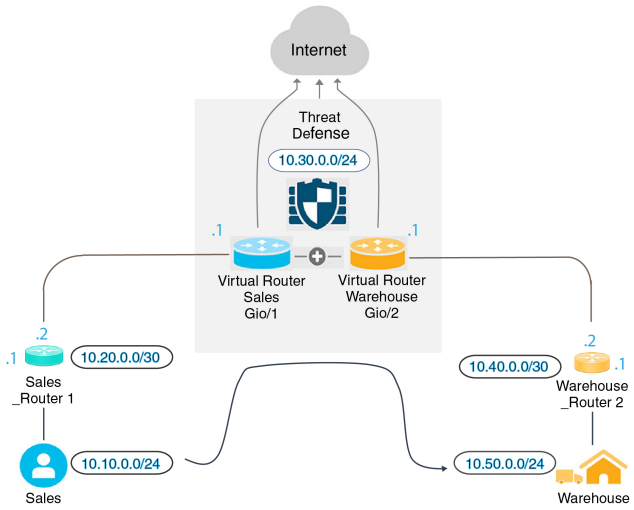

Before you begin

This example assumes that you have already configured Sales_Router1 to route traffic from 10.20.0.1/30 interface to 10.50.0.5/24.

Procedure

Step 1 Configure the inside interface (Gi0/1) of the device to be assigned to Sales virtual router:

a) Choose **Devices** > **Device Management**, click the device name, and click **Interfaces**.

b) Edit the Gi0/1 interface:

- **Name**—For this example, VR-Sales.
- Select the **Enabled** checkbox.
- In **IPV4**, for **IP Type**, choose **Use Static IP**.
- **IP Address**—Enter 10.30.0.1/24.

c) Click **Ok**.

d) Click **Save**.

Step 2 Configure the inside interface (Gi0/2) of the device to be assigned to Warehouse virtual router:

a) Choose **Devices** > **Device Management**, click the device name, and click **Interfaces**.

b) Edit the Gi0/2 interface:

- **Name**—For this example, VR-Warehouse.
- Select the **Enabled** checkbox.
- In **IPV4**, for **IP Type**, choose **Use Static IP**.
- **IP Address**—Leave it blank. The system does not allow you to configure interfaces with same IP address (10.30.0.1/24), as you are yet to create user-defined virtual routers.

c) Click **Ok**.
d) Click **Save**.

Step 3 Create Sales and Warehouse virtual routers and assign their interfaces:

a) Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
b) Choose **Routing** > **Manage Virtual Routers**.
c) Click **Add Virtual Router** and create Sales.
d) Click **Add Virtual Router** and create Warehouse.
e) Select Sales from virtual router drop-down, in **Virtual Router Properties**, add VR-Sales as **Selected Interface** and save.
f) Select Warehouse from virtual router drop-down, in **Virtual Router Properties**, add VR-Warehouse as **Selected Interface** and save.

Step 4 Revisit the VR-Warehouse interface configuration:

a) Choose **Devices** > **Device Management**, click the device name, and click **Interfaces**.
b) Click **Edit** against VR-Warehouse interface. Specify the IP Address as 10.30.0.1/24. The system now allows you to configure with same IP address of VR-Sales, because the interfaces are seperately assigned to two different virtual routers.
c) Click **Ok**.
d) Click **Save**.

Step 5 Create network objects for the warehouse server—10.50.0.0/24, and for the warehouse gateway— 10.40.0.2/30:

a) Choose **Objects** > **Object Management**.
b) Choose **Add Network** > **Add Object**:

- **Name**—For this example, Warehouse-Server.
- **Network**—Click Network and enter 10.50.0.0/24.

c) Click **Save**.
d) Choose **Add Network** > **Add Object**:

- **Name**—For this example, Warehouse-Gateway.
- **Network**—Click Host and enter 10.40.0.2.

e) Click **Save**.

Step 6 Define the route leak in Sales that points to the VR-Warehouse interface:

a) Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
b) Choose **Routing**.
c) Choose Sales virtual router from the drop-down, and then click **Static Route**.
d) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

- **Interface**—Select VR-Warehouse.
- **Network**—Select the Warehouse-Server object.
- **Gateway**—Leave it blank. When leaking a route into another virtual router, do not select the gateway.

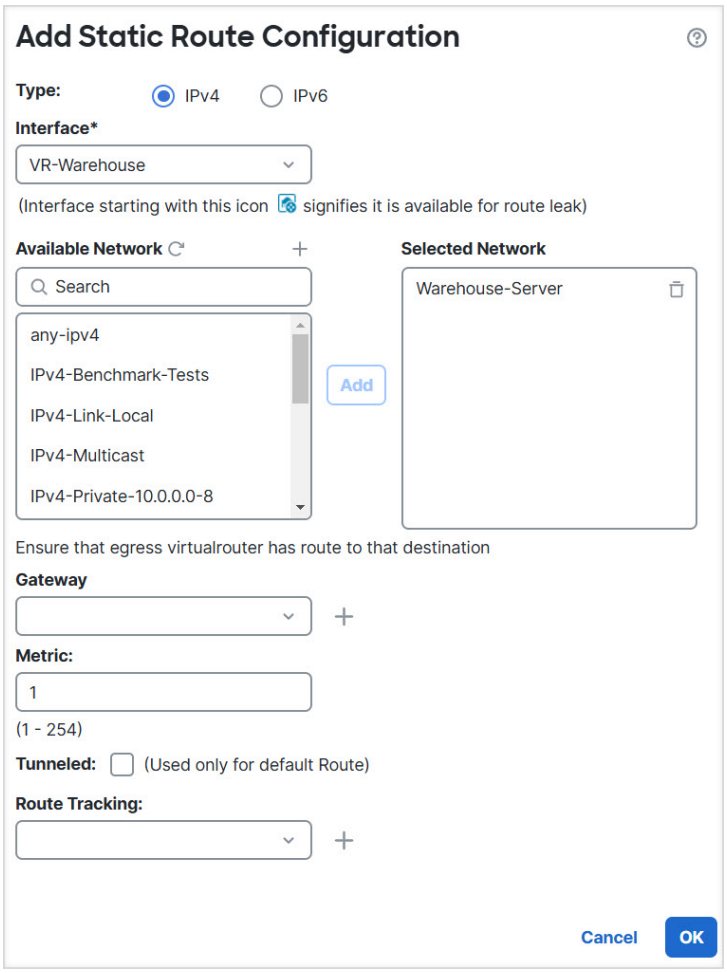

e) Click **Ok**.
f) Click **Save**.

Step 7 In the Warehouse virtual router, define the route that points to the Warehouse Router 2 gateway:

a) Choose Warehouse virtual router from the drop-down, and then click **Static Route**.
b) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

- **Interface**—Select VR-Warehouse.
- **Network**—Select the Warehouse-Server object.
- **Gateway**—Select the Warehouse-Gateway object.

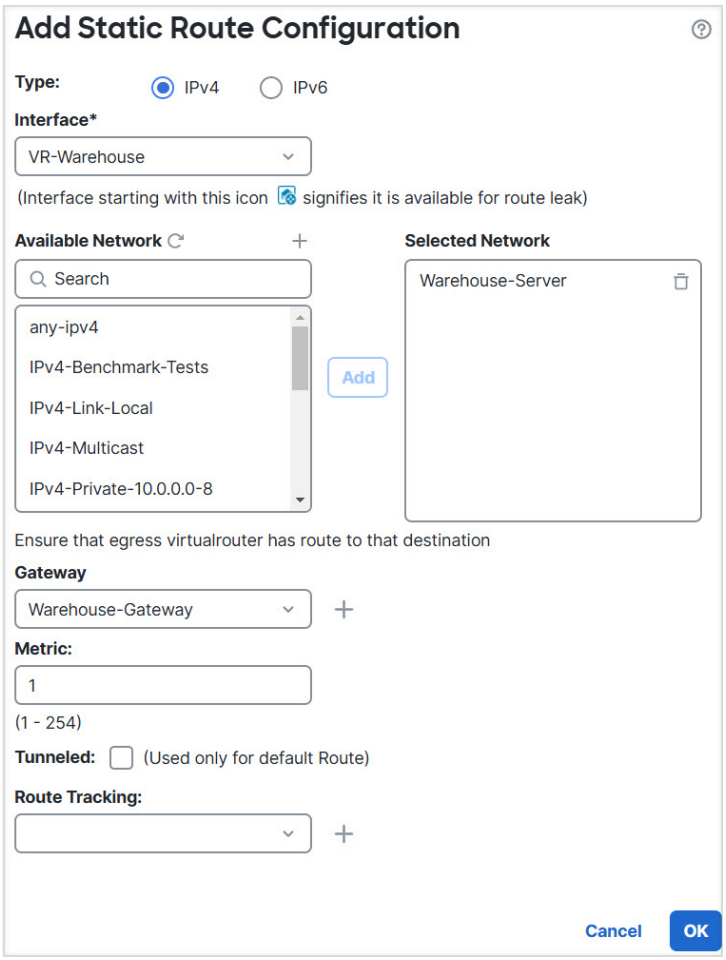

c) Click **Ok**.

d) Click **Save**.

Step 8 Configure access control rule that allows access to the warehouse server. For creating the access control rule, you need to create security zones. Use **Objects** > **Object Management** > **Interface**. Choose **Add** > **Security Zone** and create security zones for VR-Sales and VR-Warehouse; for Warehouse-Server network object, create a Warehouse-Server interface group (Choose **Add** > **Interface Group**).

Step 9 Choose **Policies** > **Access Control heading** > **Access Control** and configure an access control rule to allow traffic from the source interfaces in the Sales virtual router to the destination interfaces in the Warehouse virtual router for the destination Warehouse-Server network object.

For example, if the interfaces in Sales are in the Sales-Zone security zone, and those in Warehouse are in the Warehouse-Zone security zone, the access control rule would look similar to the following:

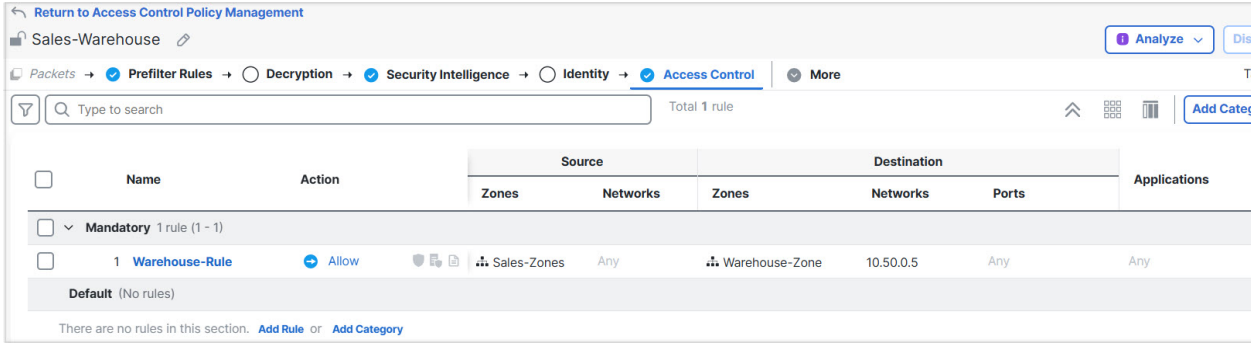

How to Provide Internet Access with Overlapping Address Spaces

When using virtual routers, you can have the same network address for interfaces that reside in separate routers. However, because the IP addresses routed in these separate virtual routers are the same, apply NAT/PAT rules for each interface with separate NAT/PAT pools to ensure that return traffic goes to the correct destination. This example provides the procedure to configure the virtual routers and NAT/PAT rules to manage the overlapping address spaces.

For example, interfaces vr1-inside and vr2-inside on Firewall Threat Defense is defined to use the IP address 192.168.1.1/24, managing endpoints on their segment in the 192.168.1.0/24 network. To allow Internet access from two virtual routers that use the same address space, you need to apply NAT rules separately to the interfaces within each virtual router, ideally using separate NAT or PAT pools. You could use PAT to translate the source addresses in VR1 to 10.100.10.1, and for those in VR2, to 10.100.10.2. The following illustration shows this setup, where the Internet-facing outside interface is part of the global router. You must define the NAT/PAT rules with the source interface (vr1-inside and vr2-inside) explicitly selected—using “any” as the source interface makes it impossible for the system to identify the correct source because the same IP address could exist on two different interfaces.

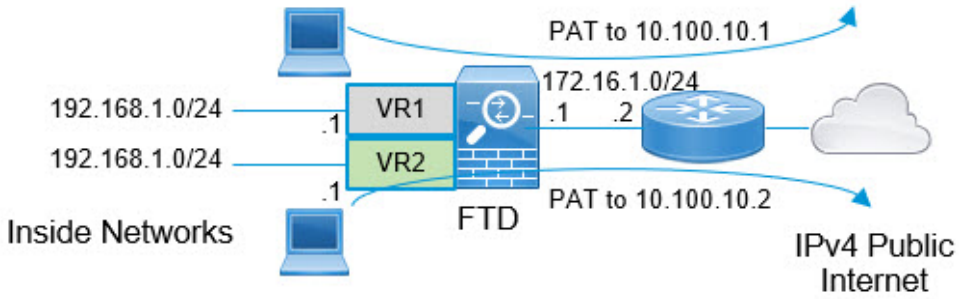

Note Even if you have some interfaces within virtual routers that does not use overlapping address spaces, define the NAT rule with the source interface to make troubleshooting easier, and to ensure a cleaner separation between traffic from the virtual routers that is Internet-bound.

Procedure

Step 1 Configure the inside interface of the device for VR1:

a) Choose **Devices** > **Device Management**, click the device name, and then click **Interfaces**.
b) Edit the interfaces that you want to assign to VR1:
   - **Name**—For this example, vr1-inside.
   - Select the **Enabled** checkbox.
   - In **IPV4**, for **IP Type**, choose **Use Static IP**.
   - **IP Address**—Enter 192.168.1.1/24.
c) Click **Ok**.
d) Click **Save**.

Step 2 Configure the inside interface of the device for VR2:

a) Choose **Devices** > **Device Management**, click the device name, and then click **Interfaces**.
b) Edit the interfaces that you want to assign to VR2:
   - **Name**—For this example, vr2-inside.
   - Select the **Enabled** checkbox.
   - In **IPV4**, for **IP Type**, choose **Use Static IP**.
   - **IP Address**—Leave it blank. The system does not allow you to configure interfaces with same IP address, as you are yet to create user-defined virtual routers.
c) Click **Ok**.
d) Click **Save**.

Step 3 Configure VR1 and the static default route leak to the outside interface:

a) Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
b) Choose **Routing** > **Manage Virtual Routers**. Click **Add Virtual Router** and create VR1.
c) For VR1, in **Virtual Router Properties**, assign vr1-inside and save.
d) Click **Static Route**.
e) Click **Add Route**. In **Add Static Route Configuration**, specify the following:
   - **Interface**—Select the outside interface of the global router.
   - **Network**—Select the any-ipv4 object. This network is the default route for any traffic that cannot be routed within VR1.
   - **Gateway**—Leave it blank. When leaking a route into another virtual router, do not provide a Gateway.

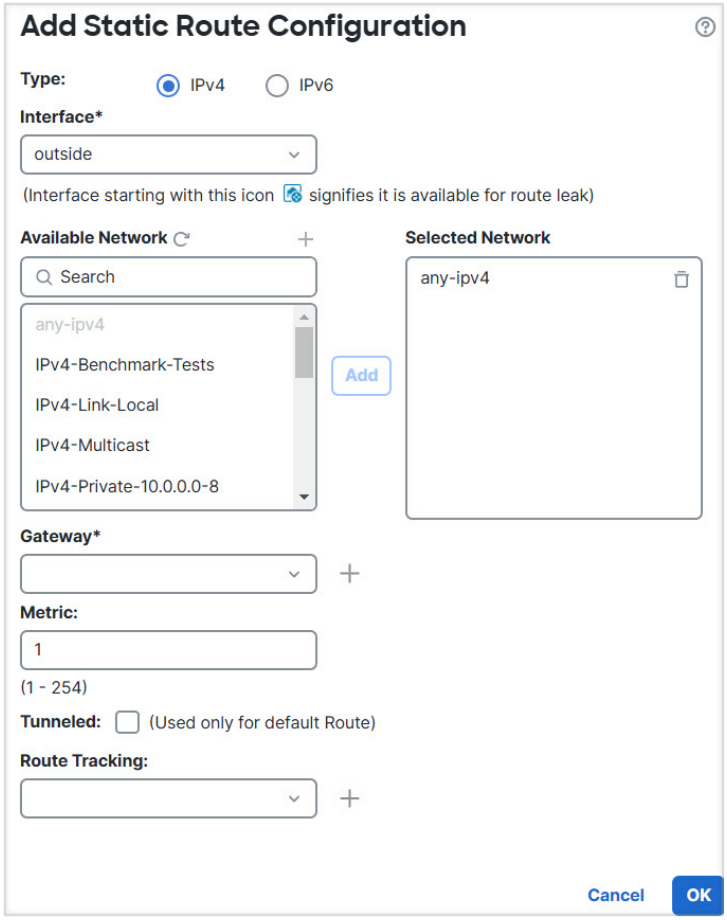

f) Click **Ok**.
g) Click **Save**.

Step 4 Configure VR2 and the static default route leak to the outside interface:

a) Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
b) Choose **Routing** > **Manage Virtual Routers**. Click **Add Virtual Router** and create VR2.
c) For VR2, in **Virtual Router Properties**, assign vr2-inside and save.
d) Click **Static Route**.
e) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

- **Interface**—Select the outside interface of the global router.
- **Network**—Select the any-ipv4 object. This network is the default route for any traffic that cannot be routed within VR2.
- **Gateway**—Leave it blank. When leaking a route into another virtual router, do not select the Gateway.

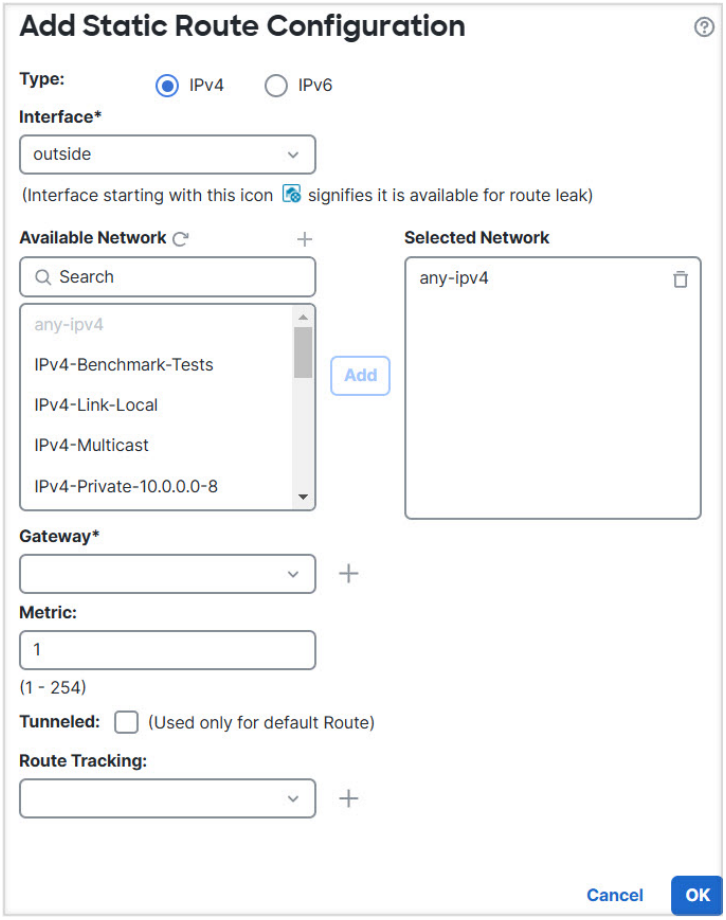

f) Click **Ok**.

g) Click **Save**.

Step 5 Configure IPv4 static default route, namely 172.16.1.2 on the outside interface of the global router:

a) Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

b) Choose **Routing** and edit global router properties.

c) Click **Static Route**.

d) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

- **Interface**—Select the outside interface of the global router.
- **Network**—Select the any-ipv4 object. This will be the default route for any IPv4 traffic.
- **Gateway**—If already created, select the host name from the drop-down. If the object is not yet created, click **Add** and define the host object for the IP address of the gateway at the other end of the network link on the outside interface, in this example, 172.16.1.2. After you create the object, select it in the Gateway field.

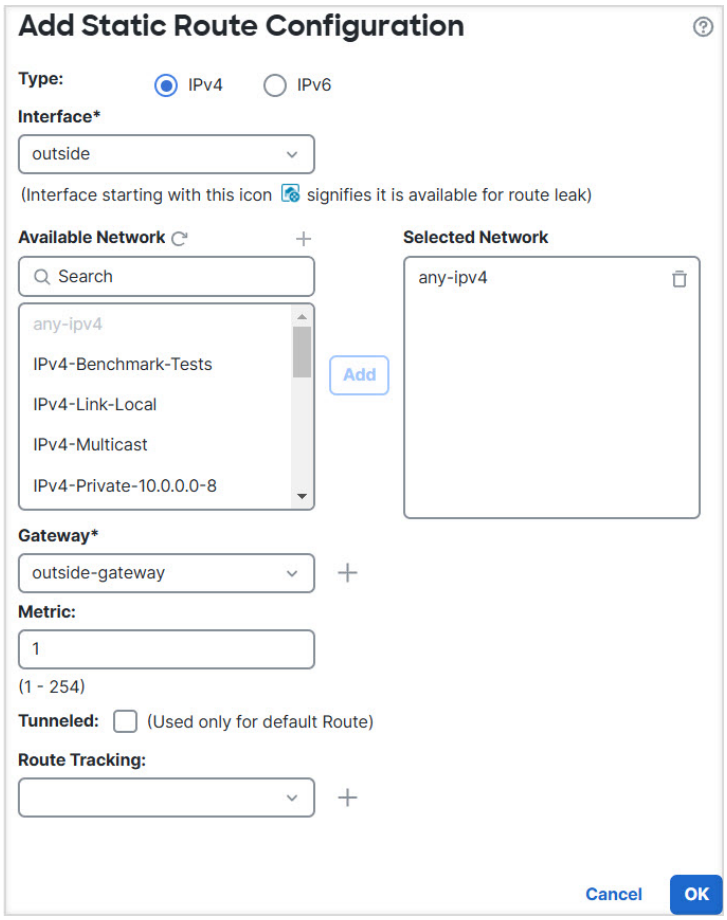

e) Click **Ok**.
f) Click **Save**.

Step 6 Revisit the vr2-inside interface configuration:

a) Choose **Devices** > **Device Management**, click the device name, and then click **Interfaces**.
b) Click **Edit** against vr2-inside interface. Specify the IP Address as 192.168.1.1/24. The system now allows you to configure with same IP address of vr1-inside, because the interfaces are separately assigned to two different virtual routers.
c) Click **Ok**.
d) Click **Save**.

Step 7 Create the NAT rule to PAT inside to outside traffic of VR1 to 10.100.10.1.

a) Choose **Devices** > **NAT** .
b) Click **New Policy**.
c) Enter InsideOutsideNATRule as the NAT policy name, and select the Firewall Threat Defense device. Click **Save**.
d) In InsideOutsideNATRule page, click **Add Rule** and define the following:

- **NAT Rule**—Select Manual NAT Rule.
- **Type**—Select Dynamic.
- **Insert**—Above, if any dynamic NAT rule exists.

- Click **Enabled**.
- In **Interface Objects**, select vr1-interface object and click **Add to Source** (If the object is not available, create one in **Objects** > **Object Management** > **Interface**), and select outside as **Add to Destination**.
- In **Translation**, for **Original Source**, select any-ipv4. For **Translated Source**, click **Add** and define host object VR1-PAT-Pool with 10.100.10.1. Select VR1-PAT-Pool as shown in the figure below:

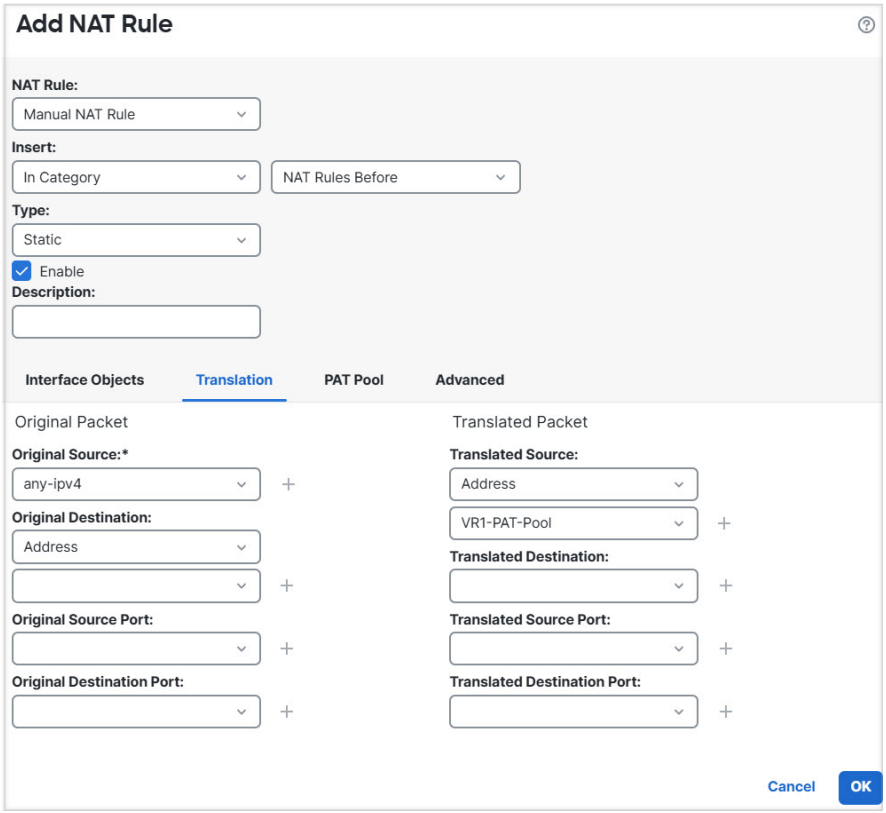

e) Click **Ok**.
f) Click **Save**.

Step 8 Add NAT rule to PAT inside to outside traffic of VR2 to 10.100.10.2.

a) Choose **Devices** > **NAT** .
b) Edit InsideOutsideNATRule to define the VR2 NAT rule:

- **NAT Rule**—Select Manual NAT Rule.
- **Type**—Select Dynamic.
- **Insert**—Above, if any dynamic NAT rule exists.
- Click **Enabled**.
- In **Interface Objects**, select vr2-interface object and click **Add to Source** (If the object is not available, create one in **System** (⚙) > **Configuration** > **Management Interfaces**), and select outside as **Add to Destination**.

- In **Translation**, for **Original Source**, select any-ipv4. For **Translated Source**, click **Add** and define host object VR2-PAT-Pool with 10.100.10.2. Select VR2-PAT-Pool as shown in the figure below:

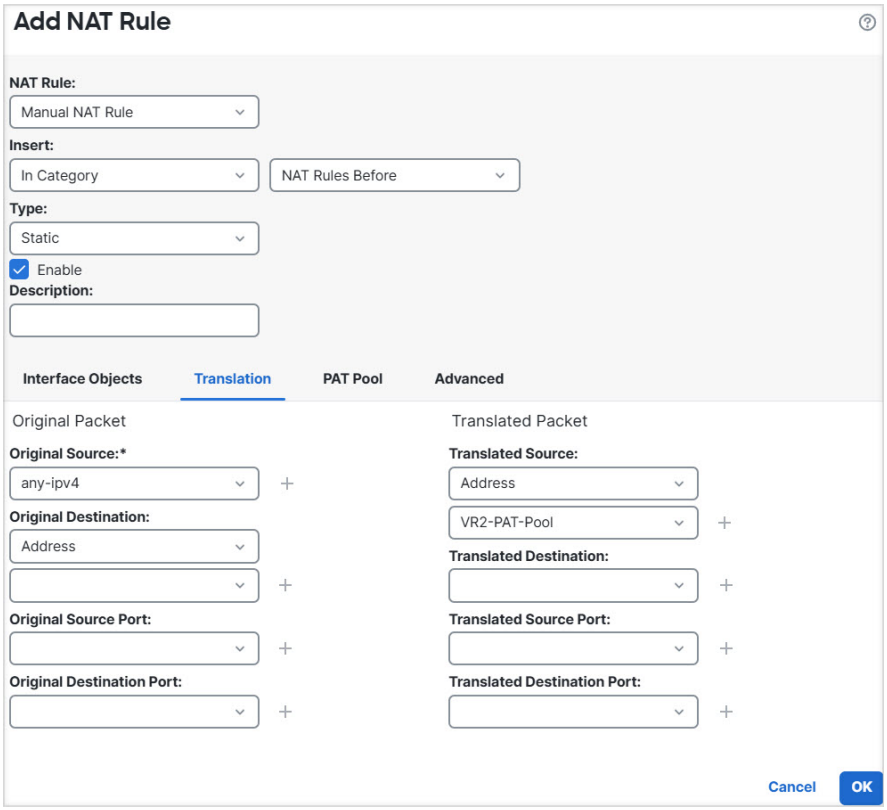

c) Click **Ok**.
d) Click **Save**.

Step 9 To configure the access control policy that allows traffic from the vr1-inside and vr2-inside interfaces to the outside interface, you need to create security zones. Use **Objects** > **Object Management** > **Interface**. Choose **Add** > **Security Zone** and create security zones for vr1-inside, vr2-inside, and outside interfaces.

Step 10 Choose **Policies** > **Access Control heading** > **Access Control** and configure an access control rule to allow traffic from vr1-inside-zone and vr2- inside-zone to outside-zone.

Assuming that you create zones named after the interfaces, a basic rule that allows all traffic to flow to the Internet will look like the following. You can apply other parameters to this access control policy:

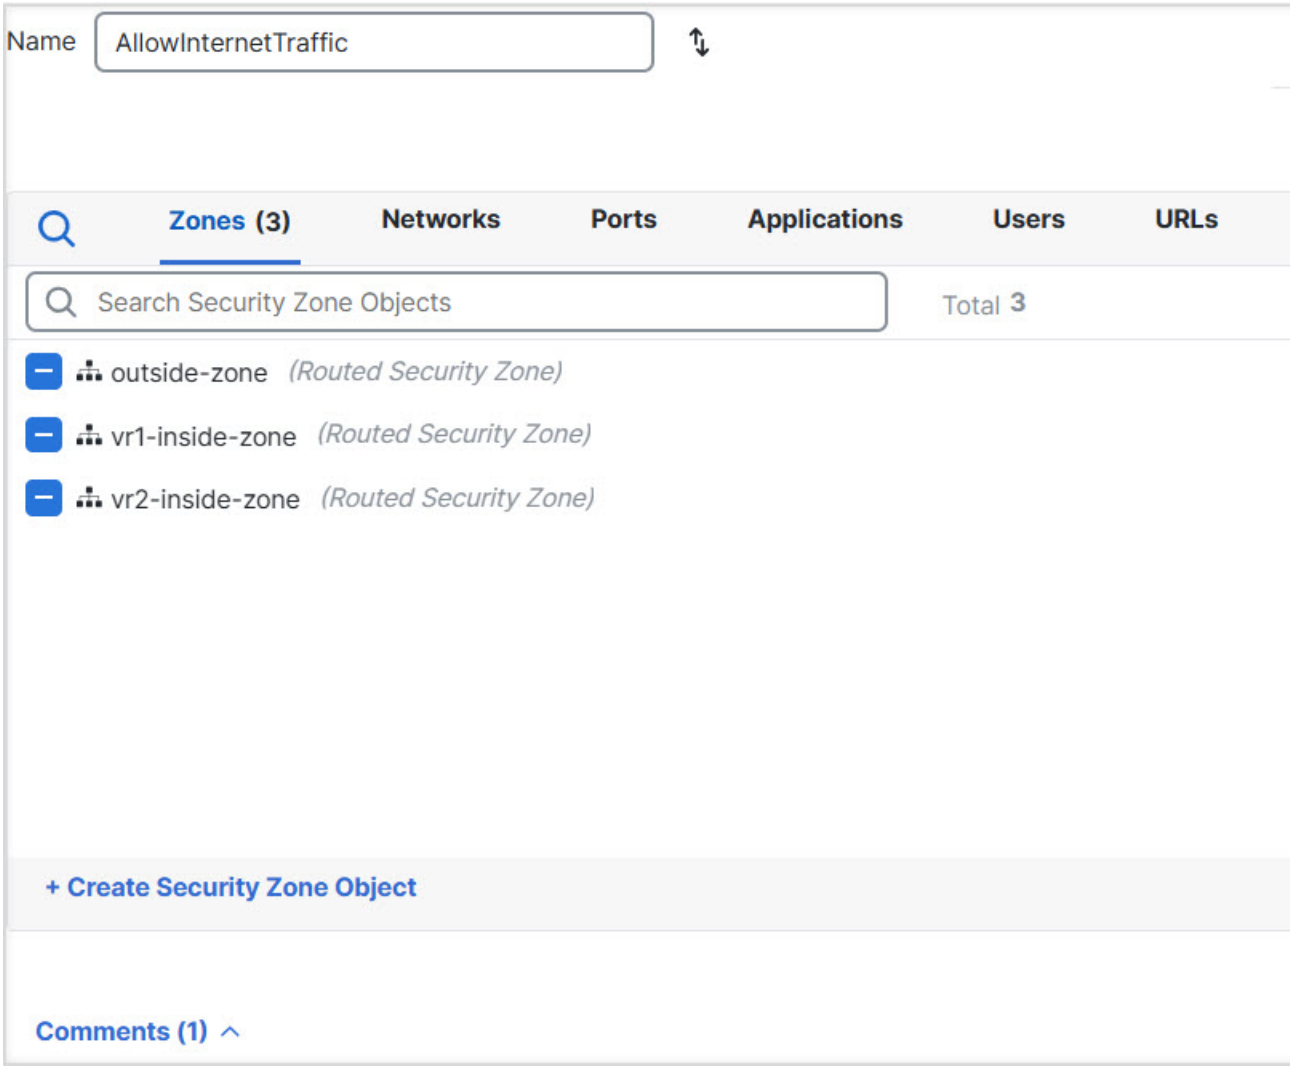

How to Allow RA VPN Access to Internal Networks in Virtual Routing

On virtual routing-enabled devices, RA VPN is supported only on global virtual router interfaces. This example provides the procedure that allows your Secure Client user to connect to user-defined virtual router networks.

In the following example, the RA VPN (Secure Client) user connects to the outside interface of Firewall Threat Defense at 172.16.3.1, and is given an IP address within the pool of 192.168.80.0/24. The user can access the inside network of only the global virtual router. To allow traffic flow through the network of the user-defined virtual router VR1, namely 192.168.1.0/24, leak the route by configuring the static routes on global and VR1.

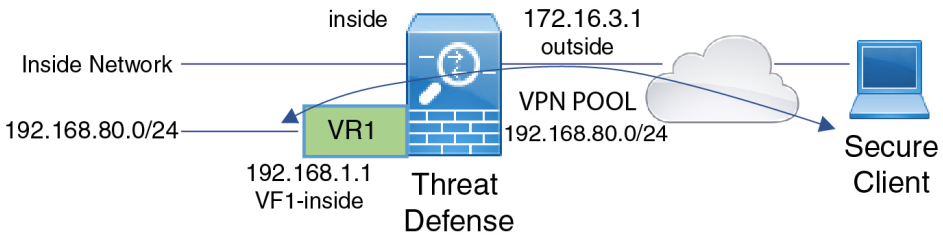

Before you begin

This example assumes that you have already configured the RA VPN, defined the virtual routers, and configured and assigned the interfaces to the appropriate virtual routers.

Procedure

Step 1 Configure route leak from Global virtual router to the user-defined VR1:

a) Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
b) Click **Routing**. By default, the Global routing properties page appears.
c) Click **Static Route**.
d) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

- **Interface**—Select the VR1 inside interface.
- **Network**—Select the VR1 virtual router network object. You can create one using the **Add Object** option.
- **Gateway**—Leave it blank. When leaking a route into another virtual router, does not select the gateway.

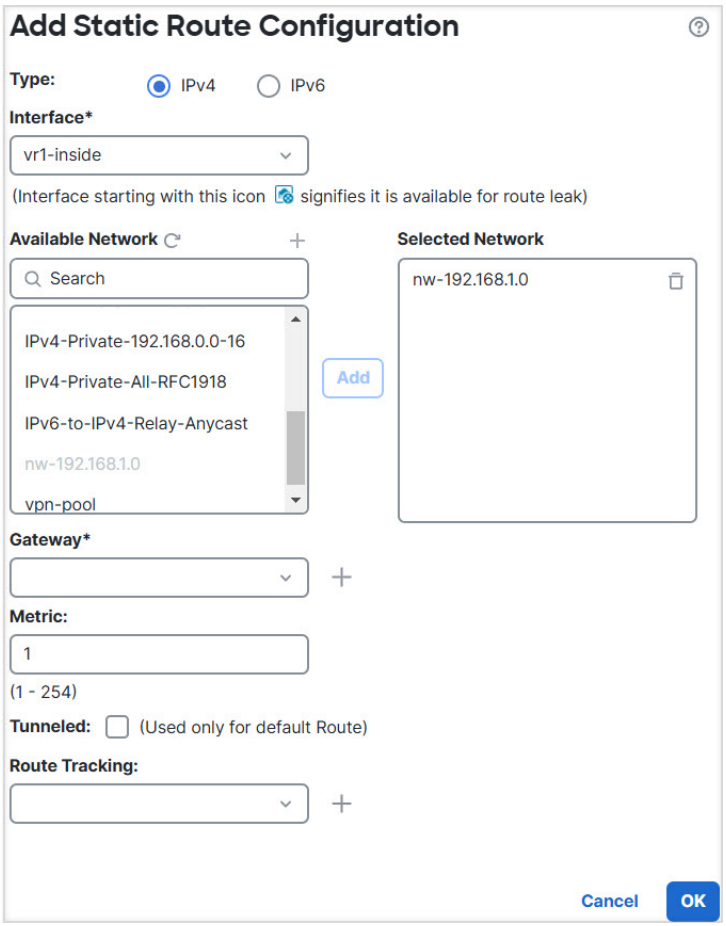

The route leak allows Secure Client assigned IP addresses in the VPN pool to access the 192.168.1.0/24 network in the VR1 virtual router.

e) Click **Ok**.

Step 2 Configure the route leak from VR1 to the Global virtual router:

a) Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
b) Click **Routing** and from the drop-down, select VR1.
c) Click **Static Route**.
d) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

- **Interface**—Select the outside interface of the global router.
- **Network**—Select the global virtual router network object.
- **Gateway**—Leave it blank. When leaking a route into another virtual router, does not select the gateway.

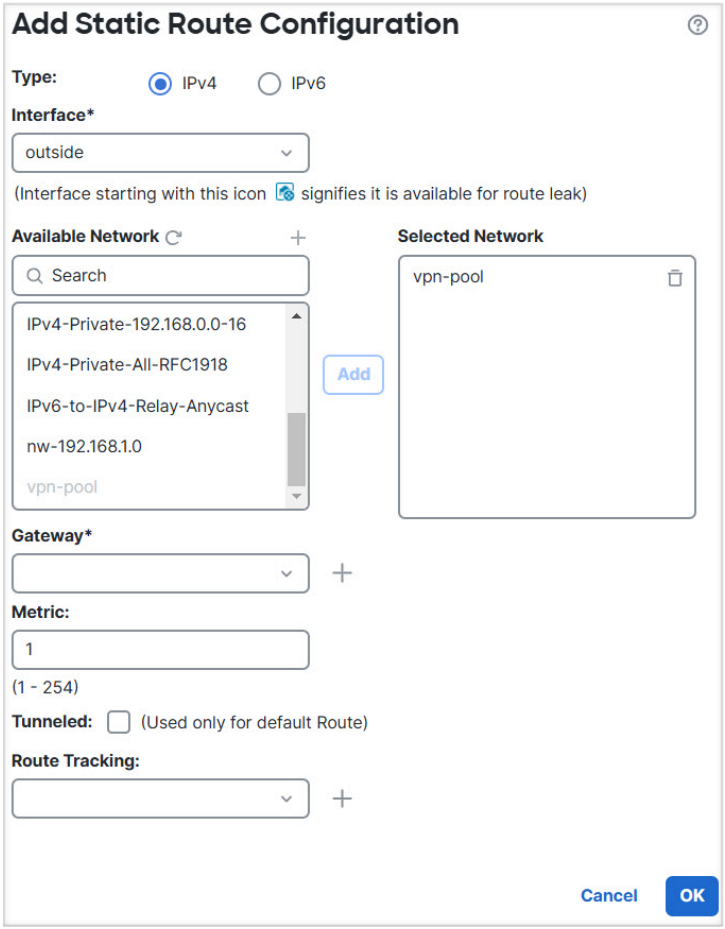

The configured static route allows endpoints on the 192.168.1.0/24 network (VR1) to initiate connections to Secure Client assigned IP addresses in the VPN pool.

e) Click **Ok**.

What to do next

If RA VPN address pool and the IP addresses in the user-defined virtual router overlap, you must also use static NAT rules on the IP addresses to enable proper routing. Alternatively, you can change your RA VPN address pool so that there is no overlap.

How to Secure Traffic from Networks in Multiple Virtual Routers over a Site-to-Site VPN

On virtual routing-enabled devices, Site-to-Site VPN is supported only on global virtual router interfaces. You cannot configure it on an interface that belongs to a user-defined virtual router. This example provides the procedure that allows you to secure the connections from or to networks hosted within user-defined virtual routers over the site-to-site VPN. You also need to update the site-to-site VPN connection to include the user-defined virtual routing networks.

Let us consider a scenario, where, a site-to-site VPN is configured between a branch office network to a company headquaters network; the Firewall Threat Defense in the branch office having virtual routers. In this case, the site-to-site VPN is defined on the outside interface of the branch office at 172.16.3.1. This VPN includes the inside network 192.168.2.0/24 without extra configuration, because the inside interface is also part of the global virtual router. But, to provide site-to-site VPN services to the 192.168.1.0/24 network, which is part of the VR1 virtual router, you must leak the route by configuring the static routes on global and VR1, and add the VR1 network to the site-to-site VPN configuration.

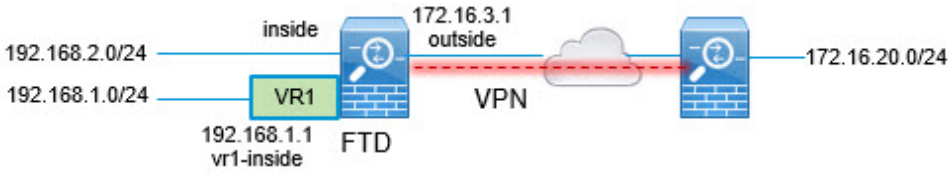

Before you begin

This example assumes that you have already configured the site-to-site VPN between the 192.168.2.0/24 local network and the 172.16.20.0/24 external network, defined the virtual routers, and configured and assigned the interfaces to the appropriate virtual routers.

Procedure

Step 1 Configure route leak from the Global virtual router to the user-defined VR1:

a) Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
b) Click **Routing**. By default, the Global routing properties page appears.
c) Click **Static Route**.
d) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

- **Interface**—Select the VR1 inside interface.
- **Network**—Select the VR1 virtual router network object. You can create one using the **Add Object** option.
- **Gateway**—Leave it blank. When leaking a route into another virtual router, do not select the gateway.

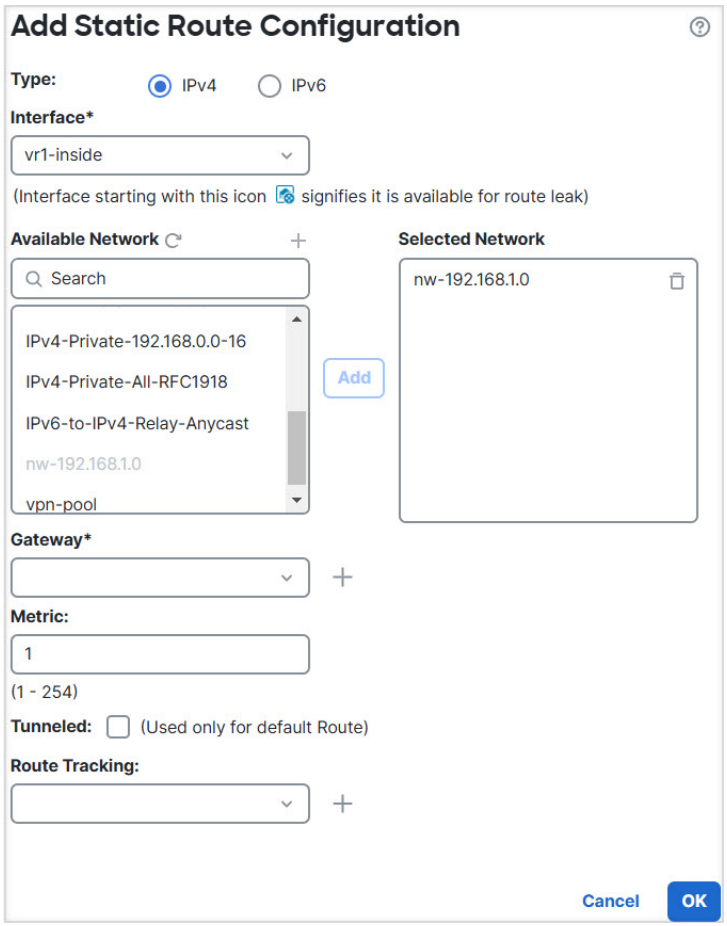

The route leak allows endpoints protected by the external (remote) end of the site-to-site VPN to access the 192.168.1.0/24 network in the VR1 virtual router.

e) Click **Ok**.

Step 2 Configure the route leak from VR1 to the Global virtual router:

a) Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
b) Click **Routing** and from the drop-down, select VR1.
c) Click **Static Route**.
d) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

- **Interface**—Select the outside interface of the global router.
- **Network**—Select the global virtual router network object.
- **Gateway**—Leave it blank. When leaking a route into another virtual router, do not select the gateway.

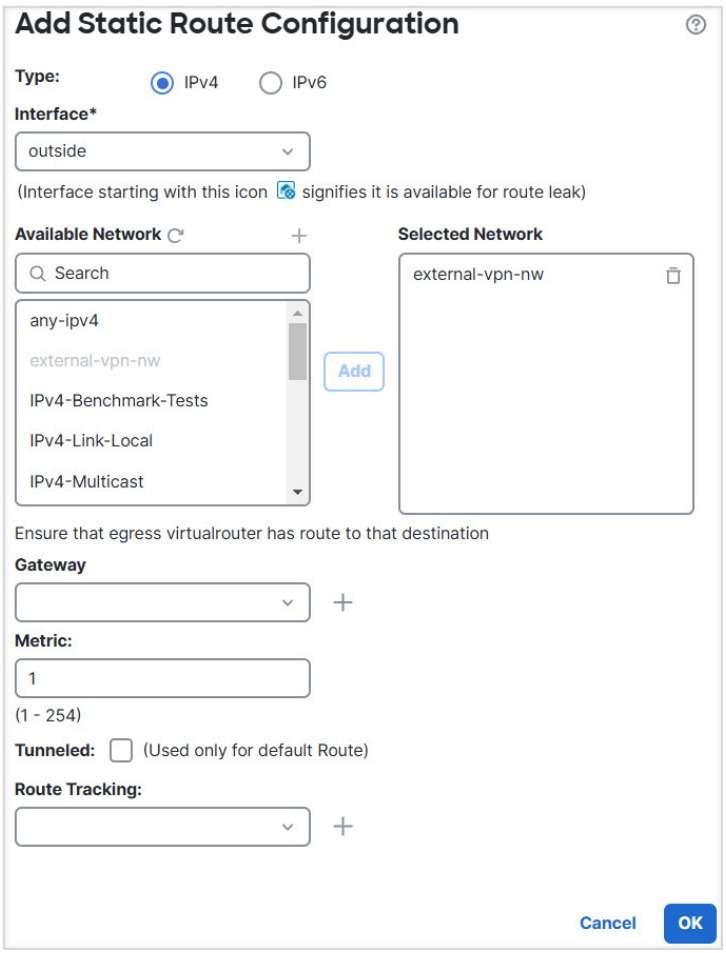

This static route allows endpoints on the 192.168.1.0/24 network (VR1) to initiate connections that will traverse the site-to-site VPN tunnel. For this example, the remote endpoint that is protecting the 172.16.20.0/24 network.

e) Click **Ok**.

Step 3 Add the 192.168.1.0/24 network to the site-to-site VPN connection profile:

a) Choose **Devices** > **VPN** > **Site to Site**, and edit the VPN Topology.
b) In **Endpoints**, edit Node A endpoint.
c) In **Edit Endpoint**, in the **Protected Networks** field, click **Add New Network Object**.
d) Add the VR1 network object with 192.168.1.0 network:

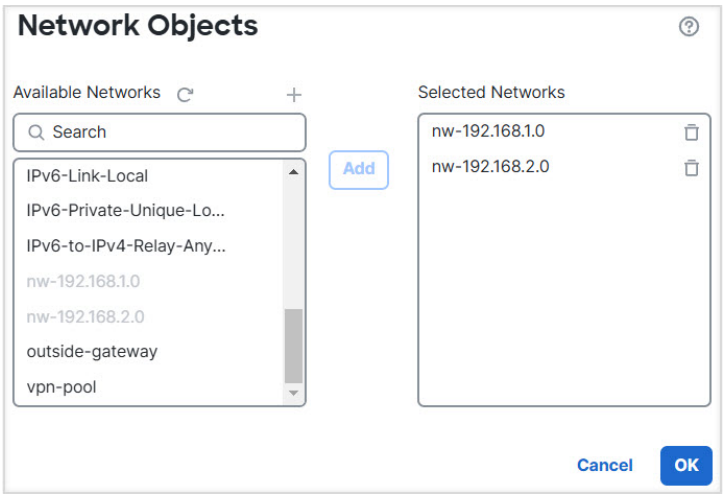

e) Click **Ok** and save the configuration.

How to Secure Traffic from Networks with Multiple Virtual Routers over a Site-to-Site VPN with Dynamic VTI

ISPs have different segmented networks for different customers. You can create virtual routers, associate dynamic VTIs with these virtual routers, and extend the capabilities of dynamic VTIs in your network. You can associate dynamic VTIs either with global or user-defined virtual routers. A single threat defense device can act as a dynamic VTI hub with a global or one or more user-defined virtual routers. Each user-defined virtual router can be one customer network.

Let us consider an example where route-based site-to-site VPNs are configured between two company headquarter networks and their two branch office networks. The ISP's threat defense, a dynamic VTI hub, manages the two company headquarter networks with two user-defined virtual routers: VRF green and VRF red. The dynamic VTI hub establishes a site-to-site VPN between:

- Customer 1 (VRF green) and Branch 1 (SVTI spoke 1)
- Customer 2 (VRF red) and Branch 2 (SVT2 spoke 2)

Figure 414: Site-to-Site VPNs with Multiple Virtual Routers and Dynamic VTI

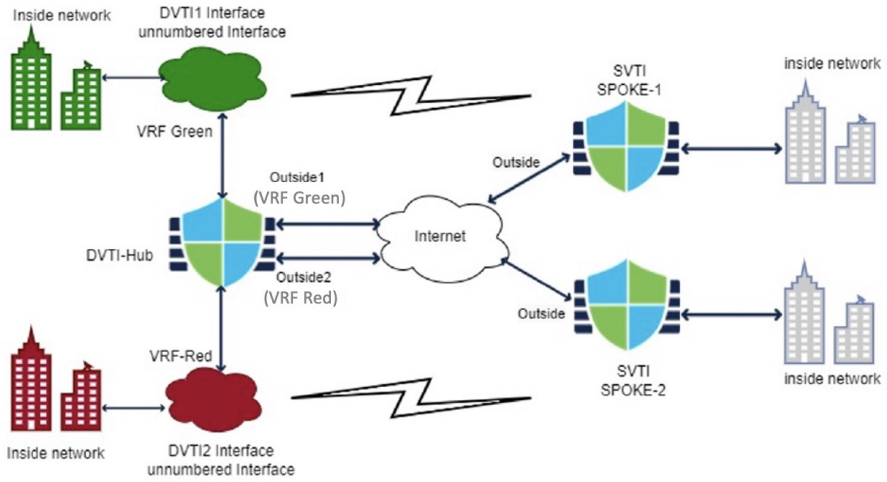

This example illustrates how to configure networks with multiple virtual routers over a site-to-site VPN with dynamic VTI.

Procedure

Step 1 Configure a dynamic VTI interface on the hub.

a) Choose **Devices** > **Device Management** and edit the threat defense device.
b) Choose **Add Interfaces > Virtual Tunnel Interface**.
c) Select the **Tunnel Type** as **Dynamic**.
d) Specify the interface name as DVTI1 and configure all the parameters for the dynamic VTI.
e) Click **Save**
f) Repeat steps 1a - e to configure the second dynamic VTI on the hub, DVTI2.

Step 2 Configure static VTI on spoke 1.

a) Choose **Devices** > **Device Management** and edit the threat defense device.
b) Choose **Add Interfaces > Virtual Tunnel Interface**.
c) Select the **Tunnel Type** as **Static**.
d) Specify the interface name as SVTI spoke-1 and configure all the parameters for the static VTI.
e) Click **Save**
f) Repeat steps 2a - e to configure the static VTI on spoke 2: SVTI spoke-2.

Step 3 Configure a route-based site-to-site VPN between hub and SVTI spoke 1.

a) Choose **Devices** > **VPN** > **Site to Site** and click **+ Site To Site VPN**.
b) Enter a name for the VPN topology in the **Topology Name** field.
c) Choose **Route Based (VTI)** and select **Hub and Spoke** as the network topology.
d) Click the **Endpoints** tab.
e) Configure the hub and the spoke (DVTI1 and SVTI spoke-1) and their routing policies.
f) Configure the **IKE**, **IPsec**, and **Advanced** options for the VPN, if required.
g) Click **Save**.

h) Repeat steps 3a - g to configure the second route-based site-to-site VPN topology between hub (DVTI2) and SVTI spoke 2.

Step 4 Configure the two virtual routers.

a) Choose **Devices** > **Device Management** and edit the threat defense device.
b) Click **Routing**.
c) Click **Manage Virtual Routers**.
d) Click **Add Virtual Router**.

Specify the name as VRF green and provide a description for the virtual router.

e) Repeat steps 4a - d to configure VRF red.

Step 5 Assign all the interfaces to the virtual routers.

a) From the drop-down list, select the virtual router.
b) In the **Virtual Router Properties** page, select the interfaces listed under the **Available Interfaces** box.

Assign the dynamic VTI interface along with the other interfaces.

c) Click **Add**.

Step 6 Repeat steps 5a - c for VRF red.

Step 7 Configure the routing policy for the virtual router.

a) From the drop-down list, select the virtual router.
b) Click **Static Route** or one of the dynamic routing protocols.
c) Configure the routing parameters.
d) Click **Save**.

What to do next

Select the hub and spoke devices and click **Deploy**. After the deployment, you can monitor the VPN tunnels in the Site-to site monitoring dashboard (**Overview** > **Dashboards** > **Site to Site VPN**).

You can also use the commands listed in Monitoring Virtual Routers, on page 1182 to view and troubleshoot the virtual router.

How to Route Traffic between Two Overlapping Network Host in Virtual Routing

You can configure hosts on the virtual routers that have same network address. If the hosts want to communicate, you can configure twice NAT. This example provides the procedure to configure the NAT rules to manage the overlapping network host.

In the following example, two hosts Host A and Host B belong to different virtual routers: VRG (interface vrg-inside), VRB (interface vrb-inside) respectively with the same subnet 10.1.1.0/24. For both the hosts to communicate, create a NAT policy where, VRG-Host interface object would use a mapped NAT address - 20.1.1.1, and VRB-Host interface object would use a mapped NAT address - 30.1.1.1. Thus, Host A uses 30.1.1.1 to communicate to Host B; Host B uses 20.1.1.1 to reach Host A.

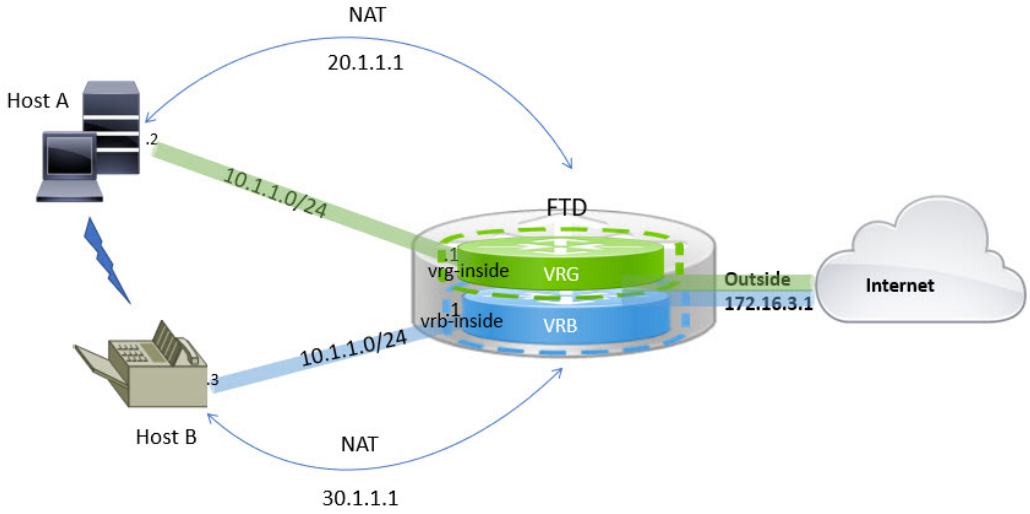

Before you begin

This example assumes that you have already configured:

- vrg-inside and vrb-inside interfaces are associated with virtual routers: VRG and VRB respectively and vrg-inside and vrb-inside interfaces configured with same subnet address (say, 10.1.1.0/24).
- Interfaces zones VRG-Inf, VRB-Inf created with vrg-inside and vrb-inside interfaces respectively.
- Host A in VRG with vrg-inside as default gateway; Host B in VRB with vrb-inside as default gateway.

Procedure

Step 1 Create the NAT rule to handle traffic from Host A to Host B. Choose **Devices** > **NAT** .

Step 2 Click **New Policy**.

Step 3 Enter a NAT policy name, and select the Firewall Threat Defense device. Click **Save**.

Step 4 In the NAT page, click **Add Rule** and define the following:

- **NAT Rule**—Select Manual NAT Rule.
- **Type**—Select Static.
- **Insert**—Select Above, if any NAT rule exists.
- Click **Enabled**.
- In **Interface Objects**, select VRG-Inf object and click **Add to Source** (If the object isn't available, create one in **Objects** > **Object Management** > **Interface**), and select VRB-Inf object and click **Add to Destination**.
- In **Translation**, select the following:
  - **Original Source**, select vrg-inside.
  - **Original Destination**, click **Add** and define object VRB-Mapped-Host with 30.1.1.1. Select VRB-Mapped-Host.

- **Translated Source**, click **Add** and define object, VRG-Mapped-Host with 20.1.1.1. Select VRG-Mapped-Host.
- **Translated Destination**, select vrb-inside as shown in the following figure:

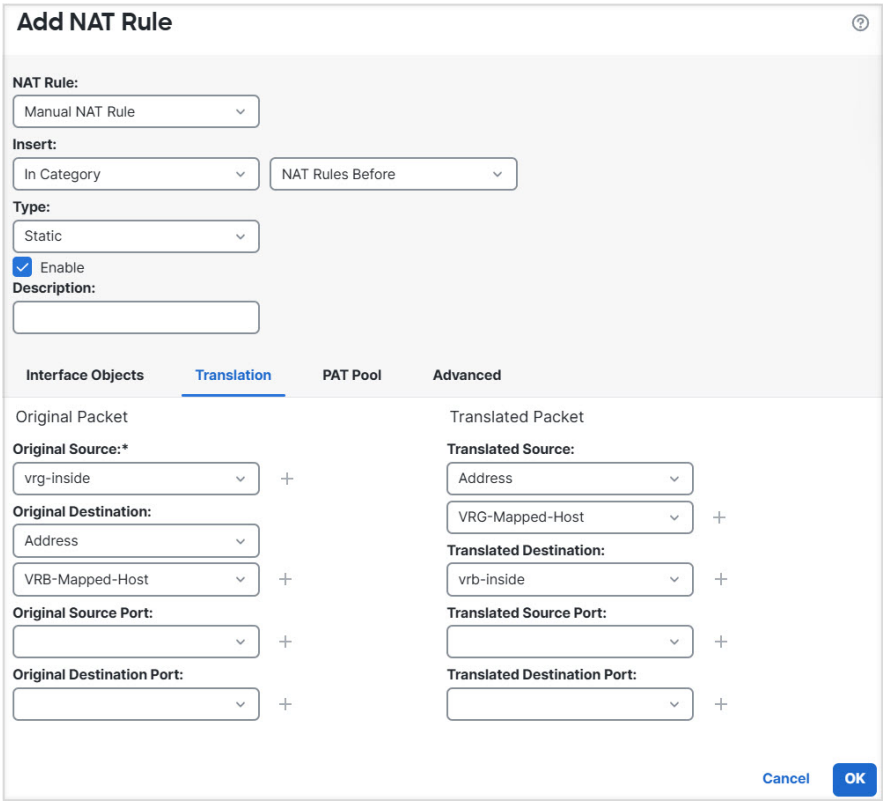

When you run the **show nat detail** command on the Firewall Threat Defense device, you will see an output similar to this:

```
firepower(config-service-object-group)# show nat detail
Manual NAT Policies (Section 1)
1 (2001) to (3001) source static vrg-inside VRG-MAPPED-HOST destination static VRB-MAPPED-HOST
 vrb-inside
translate_hits = 0, untranslate_hits = 0
Source - Origin: 10.1.1.1/24, Translated: 20.1.1.1/24
Destination - Origin: 30.1.1.1/24, Translated: 10.1.1.1/24
```

Step 5 Click **Ok**.

Step 6 Click **Save**.

The NAT rule looks like this:

When you deploy the configuration, a warning message appears:

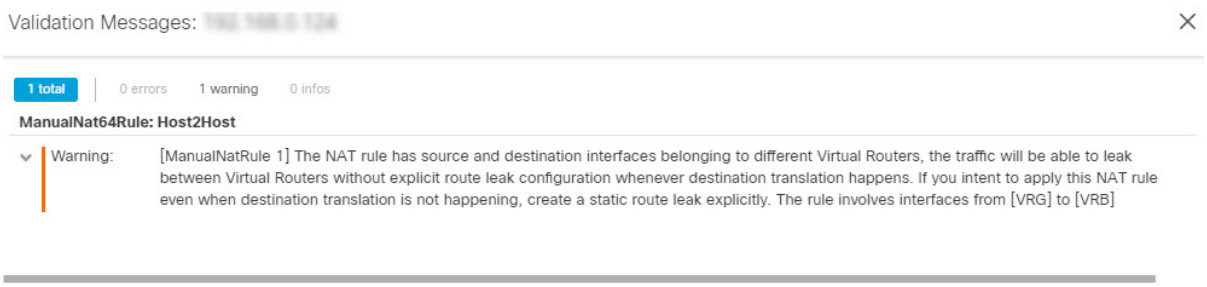

How to Manage Overlapping Segments in Routed Firewall Mode with BVI Interfaces

You can deploy single Firewall Threat Defense between multiple overlapping networks transparently and/or deploy the firewall between the hosts of same network. To achieve this deployment, configure BVI per virtual router. The procedure to configure the BVIs in virtual router is explained here.

BVI is a virtual interface within a router that acts like a normal routed interface. It does not support bridging, but represents the comparable bridge group to routed interfaces within the router. All the packets coming in or going out of these bridged interfaces, pass through the BVI interface. The interface number of the BVI is the number of the bridge group that the virtual interface represents.

In the following example, BVI-G is configured in VRG and Bridge Group 1 is the routed interface for interfaces G0/1 and G0/2. Similarly, BVI-B is configured in VRB and Bridge Group 2 is the routed interface for interfaces G0/3 and G0/4. Consider that both BVIs have the same IP subnet address, say 10.10.10.5/24. Because of virtual routers, the network is isolated on the shared resources.

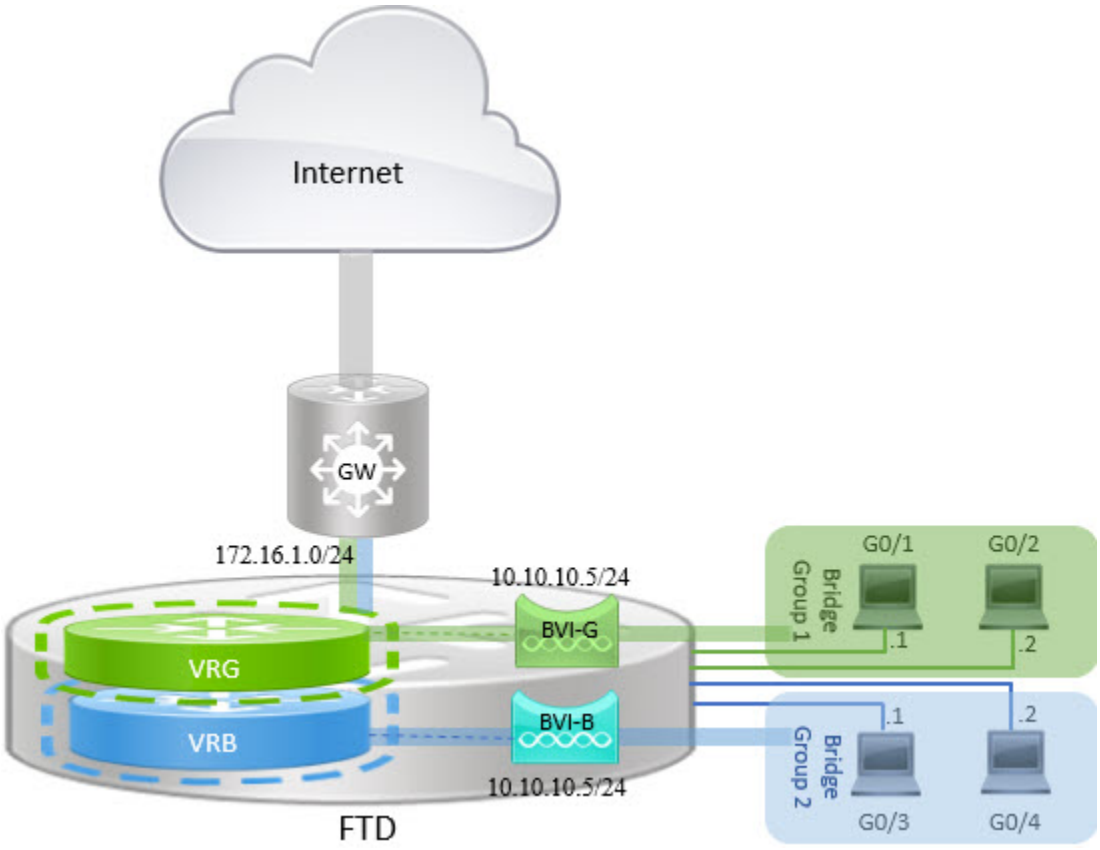

Procedure

Step 1 Choose **Devices** > **Device Management**. Edit the required device.

Step 2 In **Interfaces**, choose **Add Interfaces** > **Bridge Group Interface**.

a) Enter the following details for BVI-G:

- **Name**—For this example, BVI-G.
- **Bridge Group ID**—For this example, 1.
- **Available Interface**—Select the interfaces.
- In **IPV4**, for **IP Type**, choose **Use Static IP**.
- **IP Address**—Enter 10.10.10.5/24.

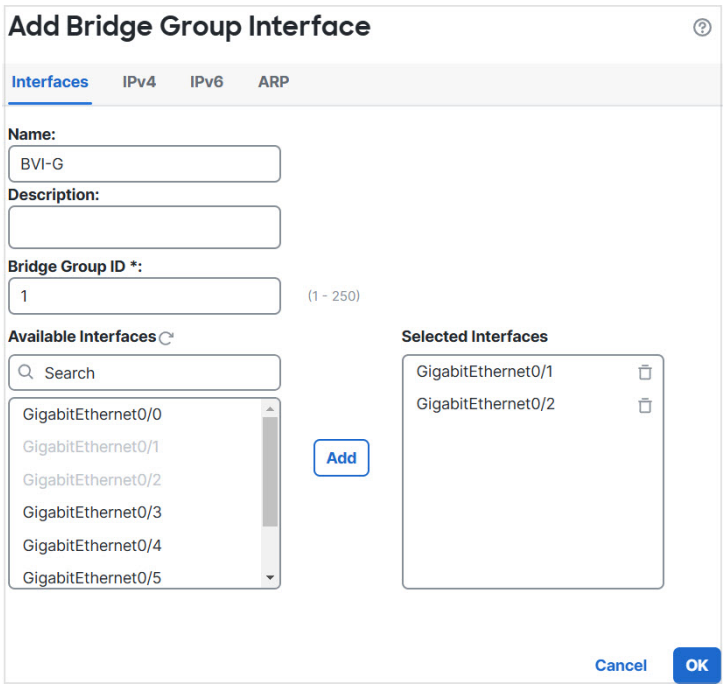

b) Click **Ok**.

c) Click **Save**.

a) Enter the following details for BVI-B:

- **Name**—For this example, BVI-B.
- **Bridge Group ID**—For this example, 2.
- **Available Interface**—Select the sub interfaces.
- In **IPV4**, for **IP Type**, choose **Use Static IP**.
- **IP Address**—Leave this field empty as the system does not allow two interfaces to have overlapping IP address. You can revisit the Bridge Group and provide the same IP address after aligning it under a virtual router.

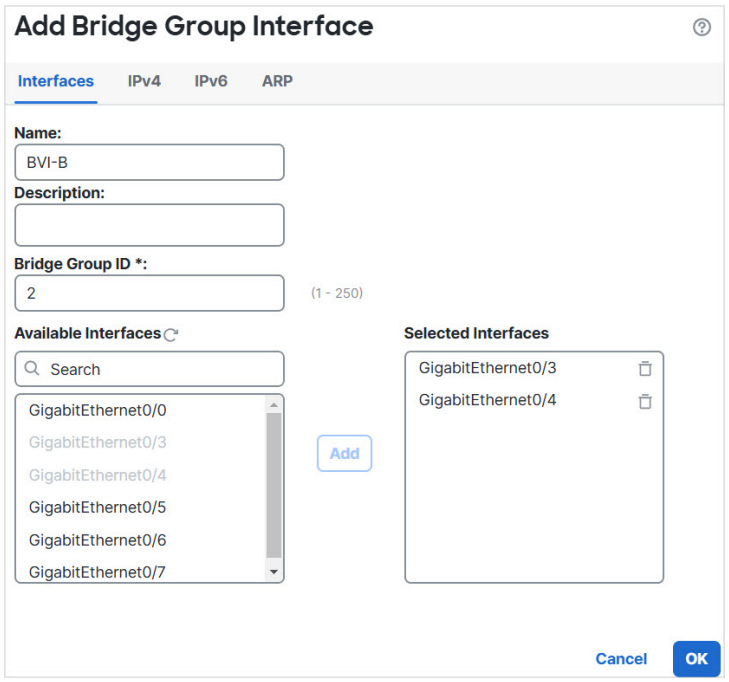

b) Click **Ok**.

c) Click **Save**.

Step 3 Create virtual router, say VRG, and select BVI-G as its network:

a) Choose **Devices** > **Device Management**.

b) Edit the device, and choose **Routing** > **Manage Virtual Routers**.

c) Click **Add Virtual Router**. Enter a name for the virtual router and click **Ok**.

d) In **Virtual Routing Properties**, select **BVI-G** and click **Add**.

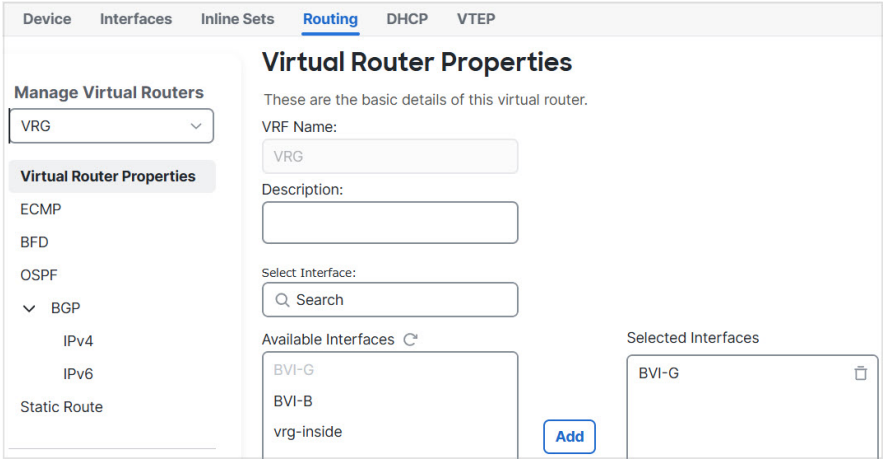

e) Click **Save**.

Step 4 Create virtual router, say VRB, and select BVI-B as its network:

a) Choose **Devices** > **Device Management**.

b) Edit the device, and choose **Routing** > **Manage Virtual Routers**.

c) Click **Add Virtual Router**. Enter a name for the virtual router and click **Ok**.
d) In **Virtual Routing Properties**, select **BVI-B** and click **Add**.

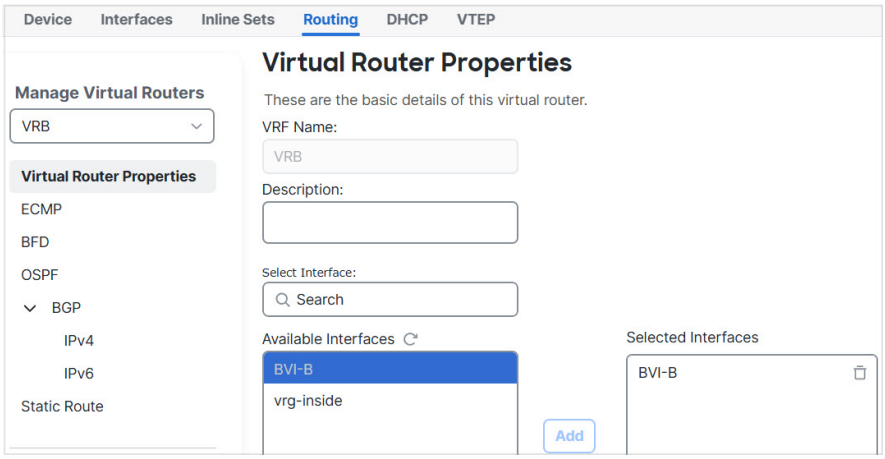

e) Click **Save**.

Step 5 Revisit the BVI-B configuration:

a) Choose **Devices** > **Device Management**, edit the device and click **Interfaces**.
b) Click **Edit** against BVI-B interface. Specify the IP Address as 10.10.10.5/24. The system now allows you to configure with same IP address of BVI-G, because the interfaces are seperately assigned to two different virtual routers.
c) Click **Ok**.
d) Click **Save**.

If you want to enable inter-BVI communication, use an external router as default gateway. In overlapping BVI scenarios, as in this example, use twice NAT external router as gateway to establish inter-BVI traffic. When configuring NAT for the members of a bridge group, you specify the member interface. You cannot configure NAT for the bridge group interface (BVI) itself. When doing NAT between bridge group member interfaces, you must specify the real and mapped addresses. You cannot specify “any” as the interface.

How to Configure User Authentication with Overlapping Networks

In virtual routing, you can configure multiple virtual routers with overlapping IP and overlapping users. In the example, VRG, and VRB are the virtual routers with overlapping IP - 192.168.1.1/24. The users on two different domains are also on overlapping network IP 192.168.1.20. For VRG and VRB users to access the shared server 172.16.10.X, leak routes to the global virtual router. Use source NAT to handle the overlapping IP. For controlling the access from VRG and VRB users, you must set user authentication in Firewall Management Center. Firewall Management Center uses realms, Active Directories, Identity source, and Identity rules and policies for authenticating user identity. Because Firewall Threat Defense does not have direct role in authenticating users, user access is managed only through the access control policy. For controlling traffic from the overlapping users, use Identity policy and rules to create access control policy.

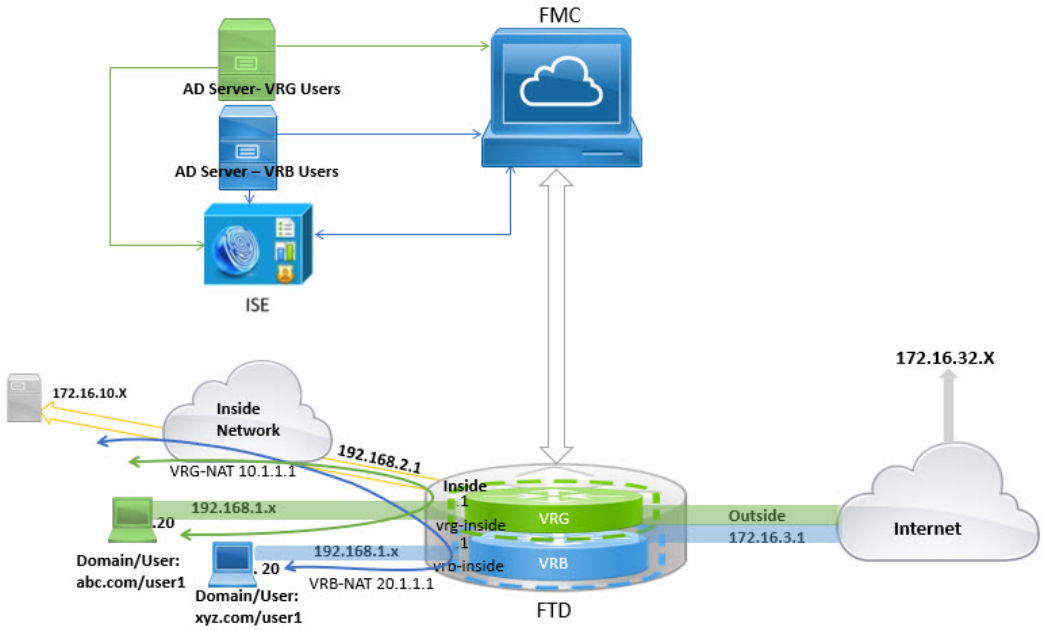

Before you begin

This example assumes that you have:

- Two AD servers for the VRG and VRB users.
- ISE with the two AD servers added.

Procedure

Step 1 Configure the inside interface of the device for VRG:

a) Choose **Devices** > **Device Management**, click the device name, and then click **Interfaces**.

b) Edit the interfaces that you want to assign to VRG:

- **Name**—For this example, VRG-inside.
- Select the **Enabled** checkbox.
- In **IPV4**, for **IP Type**, choose **Use Static IP**.
- **IP Address**—Enter 192.168.1.1/24.

c) Click **Ok**.

d) Click **Save**.

Step 2 Configure the inside interface of the device for VRB:

a) Choose **Devices** > **Device Management** > **Interfaces**.

b) Edit the interfaces that you want to assign to VRB:

- **Name**—For this example, VRB-inside.

- Select the **Enabled** checkbox.
- In **IPV4**, for **IP Type**, choose **Use Static IP**.
- **IP Address**—Leave it blank. The system doesn't allow you to configure interfaces with same IP address, as you're yet to create user-defined virtual routers.

c) Click **Ok**.
d) Click **Save**.

Step 3 Configure VRG and the static default route leak to the inside interface of the Global router for the VRG users to access the common server 172.16.10.1:

a) Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
b) Choose **Routing** > **Manage Virtual Routers**. Click **Add Virtual Router** and create VRG.
c) For VRG, in **Virtual Router Properties**, assign VRG-inside and save.

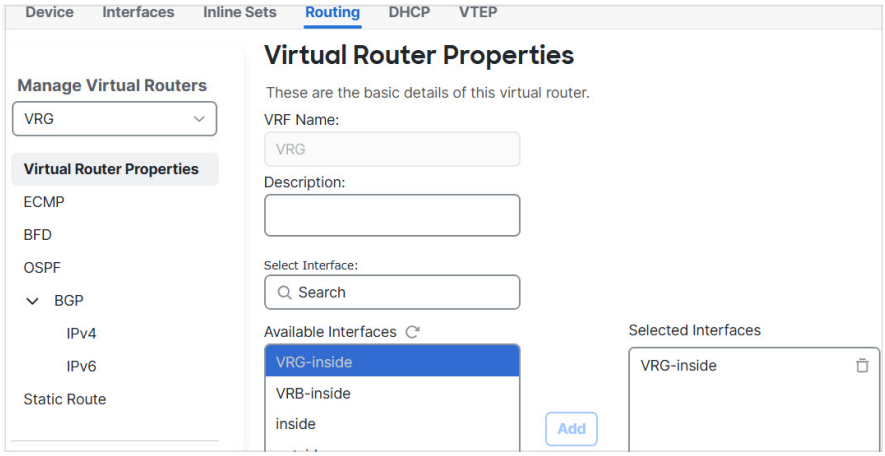

d) Click **Static Route**.
e) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

- **Interface**—Select the inside interface of the global router.
- **Network**—Select the any-ipv4 object.
- **Gateway**—Leave it blank. When leaking a route into another virtual router, do not select a gateway.

f) Click **Ok**.
g) Click **Save**.

Step 4 Configure VRB and the static default route leak to the inside interface of the Global router for the VRB users to access the shared server 172.16.10.x:

a) Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
b) Choose **Routing** > **Manage Virtual Routers**. Click **Add Virtual Router** and create VRB.
c) For VRB, in **Virtual Router Properties**, assign VRB-inside and save.

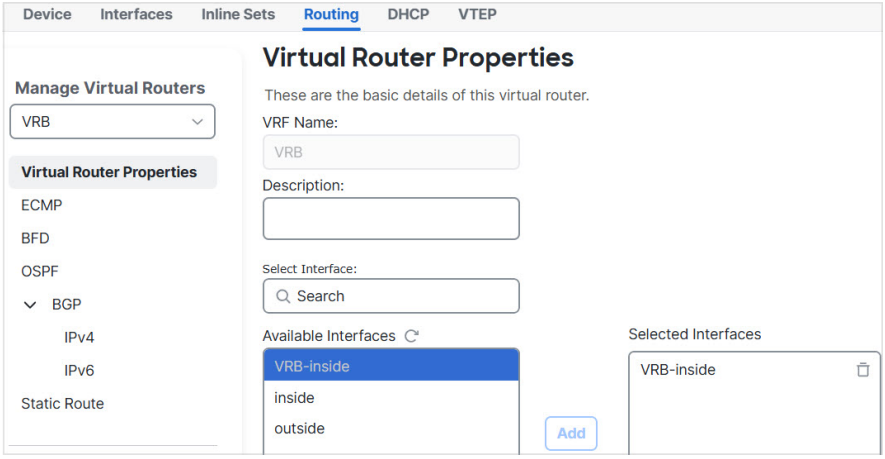

d) Click **Static Route**.
e) Click **Add Route**. In **Add Static Route Configuration**, specify the following:

- **Interface**—Select the inside interface of the global router.
- **Network**—Select the any-ipv4 object.
- **Gateway**—Leave it blank. When leaking a route into another virtual router, do not select a gateway.

f) Click **Ok**.
g) Click **Save**.

Step 5 Revisit the VRB-inside interface configuration:

a) Choose **Devices** > **Device Management** > **Interfaces**.
b) Click **Edit** against VRB-inside interface. Specify the IP Address as 192.168.1.1/24. The system now allows you to configure with the same IP address as that of VRG-inside, because the interfaces are seperately assigned to two different virtual routers.
c) Click **Ok**.
d) Click **Save**.

Step 6 Add NAT rules for the source objects VRG and VRB. Click **Devices** > **NAT**.

Step 7 Click **New Policy**.

Step 8 Enter a NAT policy name, and select the Firewall Threat Defense device. Click **Save**.

Step 9 In the NAT page, click **Add Rule** and define the following source NAT for VRG:

- **NAT Rule**—Select Manual NAT Rule.
- **Type**—Select Static.
- **Insert**—Select Above, if any NAT rule exists.
- Click **Enabled**.
- In **Interface Objects**, select VRG-Inside object and click **Add to Source** (If the object is not available, create one in **Object** > **Object Management** > **Interface**), and select Global-Inside object and click **Add to Destination**.
- In **Translation**, select the following:
  - **Original Source**, select VRG-Users.

- **Translated Source**, click **Add** and define object, VRG-NAT with 10.1.1.1. Select VRG-NAT as shown in the following figure:

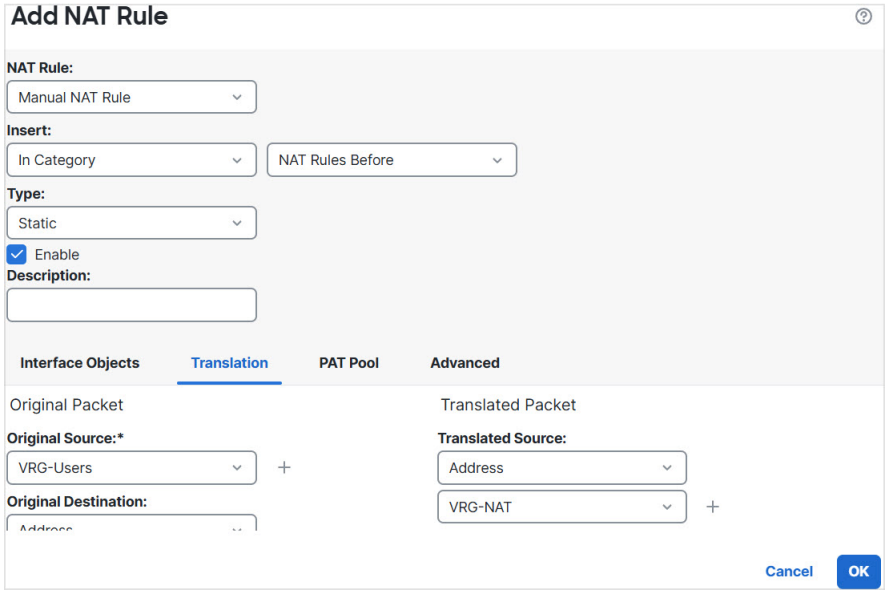

Step 10 Click **Ok**.

Step 11 In the NAT page, click **Add Rule** and define the following source NAT for VRB:

- **NAT Rule**—Select Manual NAT Rule.
- **Type**—Select Static.
- **Insert**—Select Above, if any NAT rule exists.
- Click **Enabled**.
- In **Interface Objects**, select VRB-Inside object and click **Add to Source** (If the object is not available, create one in **Object** > **Object Management** > **Interface**), and select Global-Inside object and click **Add to Destination**.
- In **Translation**, select the following:
  - **Original Source**, select VRB-Users.
  - **Translated Source**, click **Add** and define object, VRB-NAT with 20.1.1.1. Select VRB-NAT as shown in the following figure:

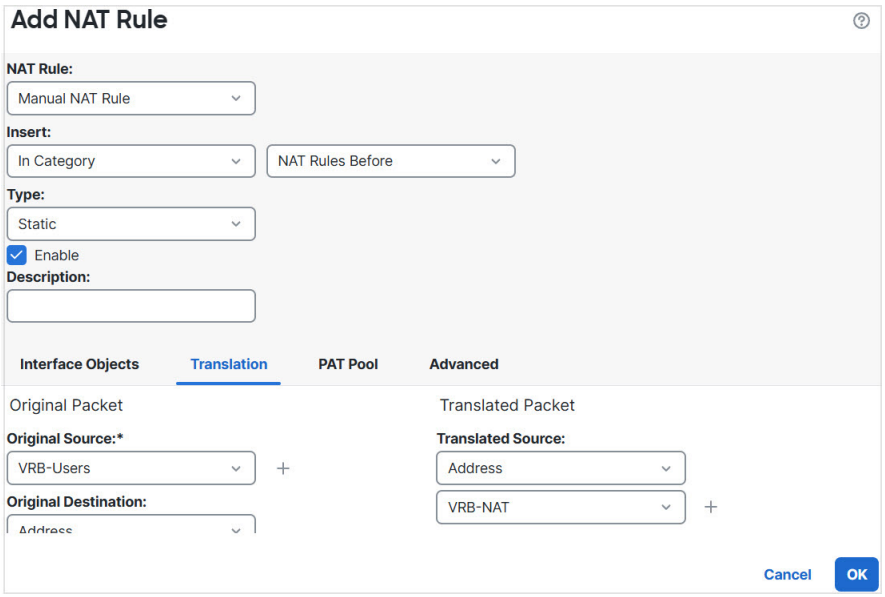

Step 12 Click **Save**.

The NAT rule looks like this:

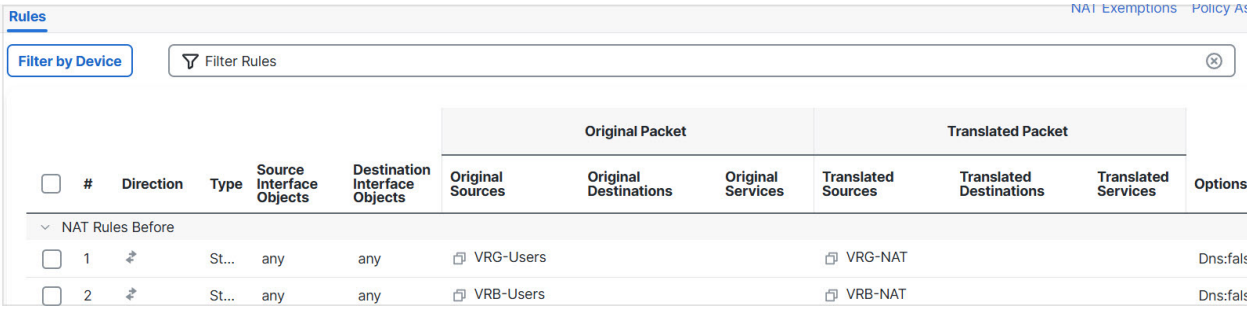

Step 13 Add the two unique AD servers in Firewall Management Center one for each VRG and VRB users—choose **System** > **Integration** > **Realms**.

Step 14 Click **New Realm** and complete the fields. For detailed information on the fields, see Realm Fields, on page 2436.

Step 15 For controlling the access from VRG and VRB users, define 2 Active Directories, see Realm Directory and Synchronize fields, on page 2440see Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426

Step 16 Add ISE in Firewall Management Center—choose **System** > **Integration** > **Identity Sources**.

Step 17 Click **Identity Services Engine** and complete the fields. For detailed information on the fields, see How to Configure ISE/ISE-PIC for User Control Using a Realm, on page 2505.

Step 18 Create Identity policy, rules, and then define access control policy for controlling access of overlapping users from VRG and VRB.

How to Interconnect Virtual Routers using BGP

You can now configure BGP settings on a device to leak the routes among virtual routers (Global and user-defined virtual routers). The route target of the source virtual router is exported to the BGP table, which, in turn is imported to the destination virtual router. The route map is used to share the Global virtual routes with the user-defined virtual routers and vice versa. Note that all import or export of the routes to the BGP table is configured at the user-defined virtual router, including the Global virtual routes.

Consider the firewall device of a factory is configured with the following virtual routers and interfaces:

- Global virtual router is configured with Inside (10.10.1.4/24) and Outside (10.10.0.5/24)
- VR-S (Sales) virtual router is configured with Inside1 (10.10.10.7/24) and Outside1 (10.10.11.7/24)
- VR-W (Warehouse) virtual router is configured with Inside2 (10.10.12.7/24) and Outside2 (10.10.13.7/24)

Assume that you want the routes of warehouse (VR-W) to be leaked with sales (VR-S) and Global, and the outside interface routes of VR-S to VR-W. Similarly, you want the outside interface routes of the Global router to be leaked to sales (VR-S). This example demonstrates the BGP configuration procedure to achieve interconnecting the routers:

Figure 415: Interconnect Virtual Routers using BGP Settings

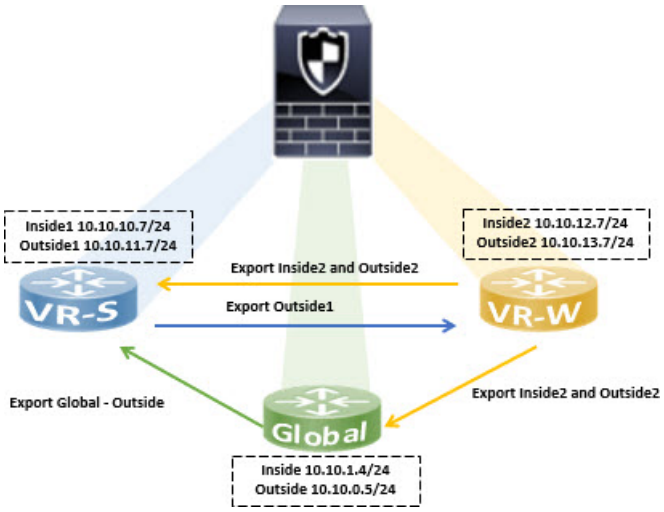

Before you begin

- Create the virtual routers—VR-S and VR-W.
- Enable BGP and for each virtual router Configure BGP for redistribution of connected routes.

Procedure

Step 1 Configure VR-W to export its routes tagging them with a route target to VR-S:

a) Choose **Devices** > **Device Management**, edit device, and then click the **Routing** tab.

b) From the virtual router drop-down, select VR-W.

c) Click **BGP** > **IPv4** > **Route Import/Export**.

d) To leak the VR-W routes to VR-S, tag the routes with a route target, so that the VR-W routes are exported to its BGP table with the route target marked on them. In the **Route Targets Export** field, enter a value, say, *200:200*. Click **Add**:

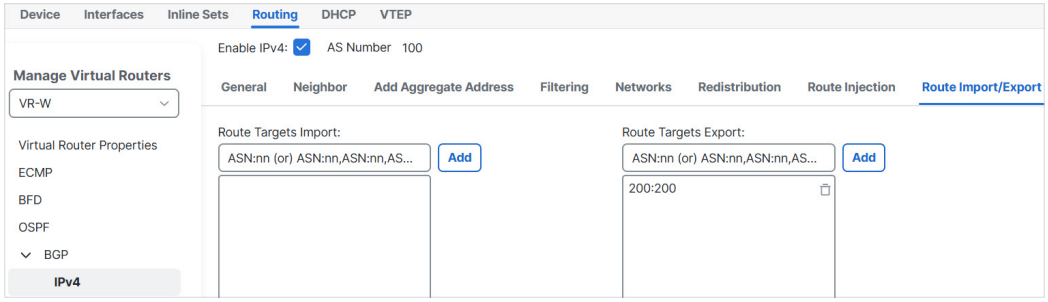

e) From the virtual router drop-down, select VR-S.

f) Click **BGP** > **IPv4** > **Route Import/Export**.

g) To receive the leaked routes from VR-W, configure the Import Route Target to import the VR-W routes that are marked with the route target from the (peer or redistributed) BGP table. In the **Route Targets Import** field, enter the same route target value that you had configured for VR-W, *200:200*. Click **Add**.

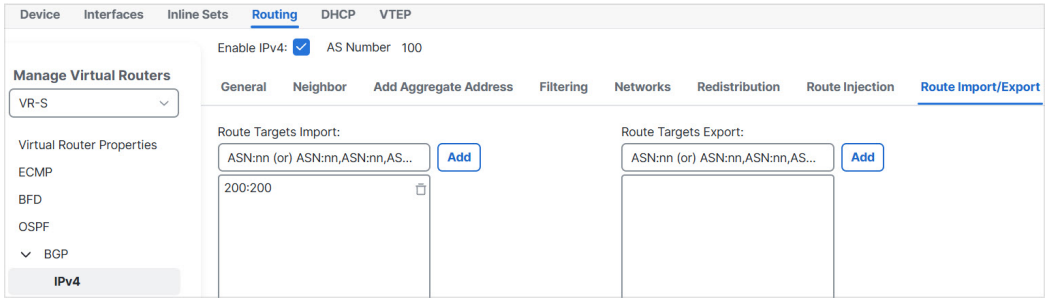

Note

If you want to conditionalize routes to be leaked from VR-W, you can specify the match criteria in the route map object, and choose it in the **User Virtual Router Export Route Map**. Similarly, if you want to conditionalize the routes to be imported to VR-S from the BGP table, you can use the **User Virtual Router Import Route Map**. This procedure is explained in Step 3.

Step 2 Configure VR-W to export its routes to the Global virtual router:

a) You need to create a route map that would allow the VR-W routes to be exported to the Global routing table. Choose **Objects** > **Object Management** > **Route Map**.

b) Click **Add Route Map**, give a name, say *Export-to-Global*, and then click **Add**.

c) Specify a **Sequence Number**, say 1, and then choose Allow from the **Redistribution** drop-down list:

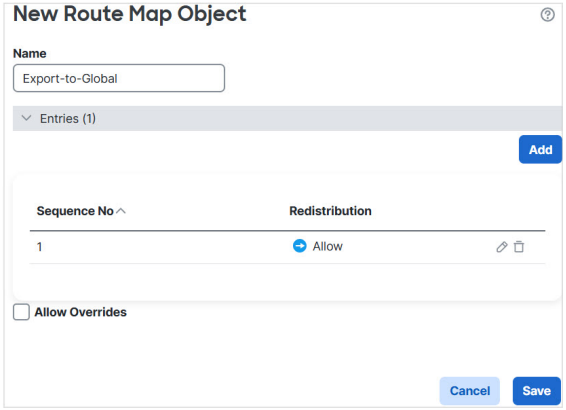

d) Click **Save**.

In this example, all the VR-W routes are leaked to the Global routing table. Hence, no match criteria is configured for the route map.

e) Navigate to the **Routing** tab of the device, and select VR-W. Click **BGP** > **IPv4** > **Route Import/Export**.

f) From the **Global Virtual Router Export Route Map** drop-down list, choose Export-to-Global:

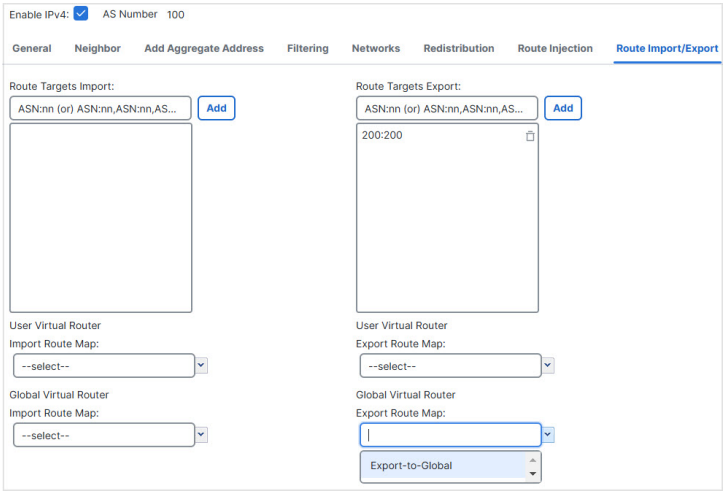

Step 3 To leak only the Outside1 routes of VR-S to VR-W:

a) From the virtual router drop-down, select VR-S.

b) Click **BGP** > **IPv4** > **Route Import/Export**.

c) To leak the VR-S routes to VR-W, tag the routes with a route target, so that the VR-S routes are exported to its BGP table with the route target marked on them. In the **Route Targets Export** field, enter a value, say, *100:100*. Click **Add**.

d) From the virtual router drop-down, select VR-W, and choose **BGP** > **IPv4** > **Route Import/Export**.

e) To receive the leaked routes from VR-S, configure the Import Route Target to import the VR-S routes that are marked with the route target from the (peer or redistributed) BGP table. In the **Route Targets Import** field, enter the VR-S route target value, *100:100*. Click **Add**.

f) Now, you need to conditionalize that only the Outside1 routes of VR-S to be leaked to VR-W. Choose **Objects** > **Object Management** > **Prefix List** > **IPv4 Prefix List**.

g) Click **Add IPv4 Prefix List**, give a name, say *VRS-Outside1-Only*, and then click **Add**.

h) Specify a **Sequence Number**, say 1, and then choose Allow from the **Redistribution** drop-down list.
i) Enter the IP Address (first two octets) of the VR-S Outside1 interface.
j) Click **Save**.
k) Create a route map with the match clause with the prefix list. Click **Route Map**. Click **Add Route Map**, give a name, say *Import-from-VRS*, and then click **Add**.
l) Specify a **Sequence Number**, say 1, and then choose Allow from the **Redistribution** drop-down list.
m) In the **Match Clause** tab, click **IPv4**. Under **Address** tab, click **Prefix List**.
n) Under **Available IPv4 Prefix List**, select VRS-Outside1-Only, and then click **Add**.
o) Click **Save**.
p) Navigate to the **Routing** tab of the device, and select VR-W. Click **BGP** > **IPv4** > **Route Import/Export**.
q) From the **Global Virtual Router Import Route Map** drop-down list, choose Import-from-VRS:

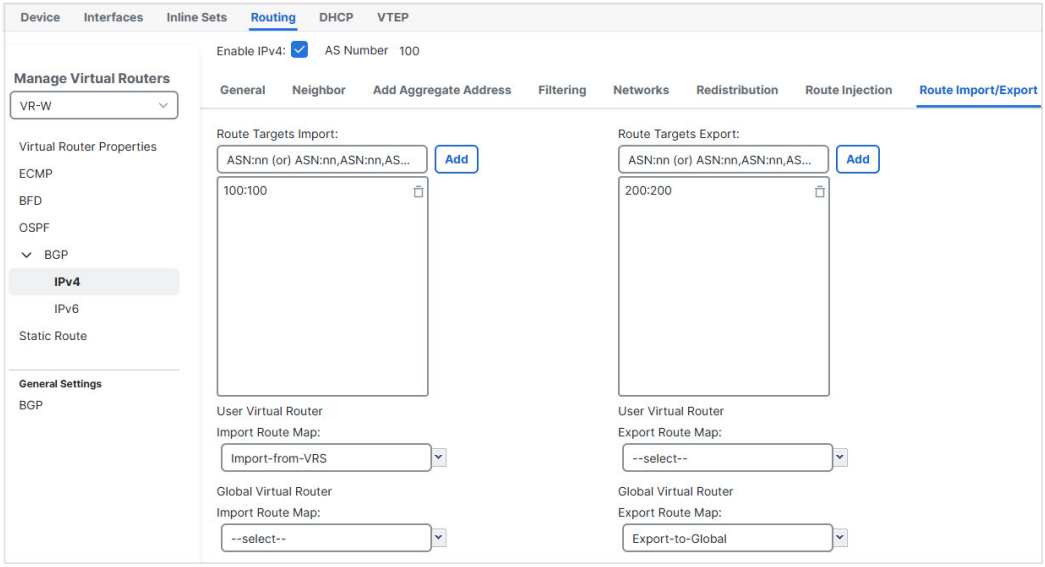

Step 4 Configure VR-S to import the Outside routes of Global virtual router:

Note
To leak routes to or from a Global virtual router, you must configure the source or destination user defined virtual router respectively. Thus, in this example, VR-S is the destination router that imports the routes from Outside interface of the Global virtual router.

a) Choose **Objects** > **Object Management** > **Prefix List** > **IPv4 Prefix List**.
b) Click **Add IPv4 Prefix List**, give a name, say *Global-Outside-Only*, and then click **Add**.
c) Specify a **Sequence Number**, say 1, and then choose Allow from the **Redistribution** drop-down list.
d) Enter the IP Address (first two octets) of the Global Outside interface:

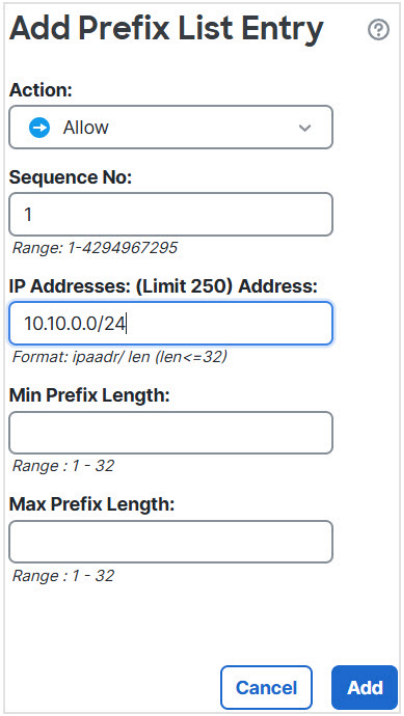

e) Click **Save**.

f) Click **Route Map**. Click **Add Route Map**, give a name, say *Import-from-Global*, and then click **Add**.

g) Specify a **Sequence Number**, say 1, and then choose Allow from the **Redistribution** drop-down list.

h) In the **Match Clause** tab, click **IPv4**. Under **Address** tab, click **Prefix List**.

i) Under **Available IPv4 Prefix List**, select Global-Outside-Only, and then click **Add**:

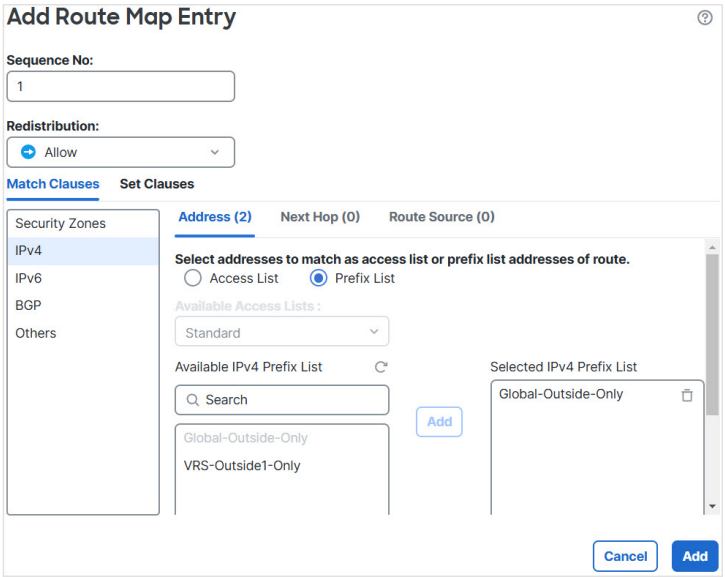

j) Click **Save**.

k) Navigate to the **Routing** tab of the device, and select VR-S. Click **BGP** > **IPv4** > **Route Import/Export**.

l) From the **Global Virtual Router Import Route Map** drop-down list, choose Import-from-Global:

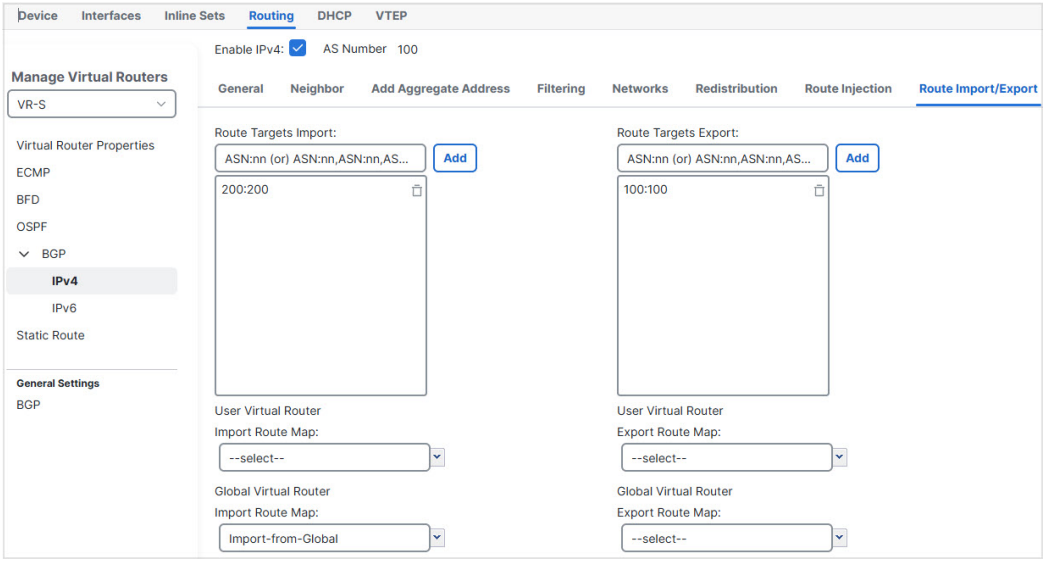

Step 5 **Save** and **Deploy**.

History for Virtual Routers

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
AAA for user-defined VRF interfaces.	7.6	7.6	A device's authentication, authorization, and accounting (AAA) is now supported on user-defined Virtual Routing and Forwarding (VRF) interfaces. The default is to use the management interface. In device platform settings, you can now associate a security zone or interface group having the VRF interface, with a configured external authentication server. New/modified screens: **Devices** > **Platform Settings** > **External Authentication**
Virtual routing with dynamic VTI.	7.4	7.4	You can now configure a virtual router with a dynamic VTI for a route-based site-to-site VPN. New/modified screens: **Devices** > **Device management** > **Edit Device** > **Routing** > **Virtual Router Properties** > **Dynamic VTI interfaces under Available Interfaces** Platform restrictions: Supported only on native mode standalone or high availability devices. Not supported for container instances or clustered devices.
Virtual router support for the ISA 3000.	7.0	7.0	You can configure up to 10 virtual routers on the ISA 3000.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
SNMP support on user-defined virtual routers	7.0	7.0	You can now configure SNMP on user-defined virtual routers.
Bulk removal of virtual routers.	6.7	6.6	You can remove multiple virtual routers from threat defense at a time. New/modified screens: **Devices** > **Device Management** > **Routing** > **Manage Virtual Routers**
Virtual routers and VRF-Lite for threat defense.	6.6	6.6	You can now create multiple virtual routers to maintain separate routing tables for groups of interfaces. Because each virtual router has its own routing table, you can provide clean separation in the traffic flowing through the device. Virtual routers implement the "light" version of Virtual Routing and Forwarding, or VRF-Lite, which does not support Multiprotocol Extensions for BGP (MBGP). The maximum number of virtual routers you can create ranges from five to 100, and depends on the device model.. New/modified screens: **Devices** > **Device Management** edit device **Routing** tab New/modified CLI commands: . • **show vrf** • Added the [**vrf** *name* | **all**] keyword set to these CLI commands, and changed the output to indicate virtual router information where applicable: **clear ospf**, **clear route**, **ping**, **show asp table routing**, **show bgp**, **show ipv6 route**, **show ospf**, **show route**, **show snort counters** Platform restrictions: Not supported on the Firepower 1010 and ISA 3000.

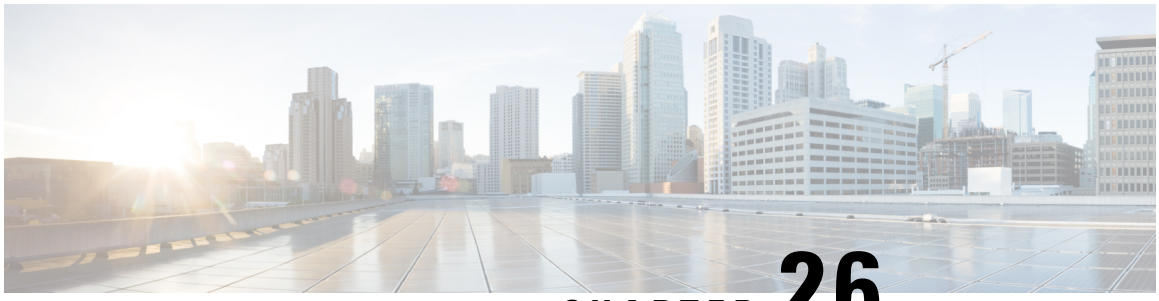

CHAPTER 26

ECMP

This chapter describes the procedure to configure Equal Cost Multi-Path (ECMP) routing that routing protocols use to load balance the network traffic.

- About ECMP, on page 1223
- Guidelines and Limitations for ECMP, on page 1223
- Manage ECMP Page, on page 1224
- Create an ECMP Zone, on page 1225
- Configure an Equal Cost Static Route, on page 1226
- Modify an ECMP Zone, on page 1227
- Remove an ECMP Zone, on page 1227
- Configuration Example for ECMP, on page 1228
- History for ECMP in Secure Firewall Threat Defense, on page 1231

About ECMP

The Firewall Threat Defense device supports Equal-Cost Multi-Path (ECMP) routing. You can configure traffic zones per virtual router to contain a group of interfaces. You can have up to 8 equal cost static or dynamic routes across up to 8 interfaces within each zone. For example, you can configure multiple default routes across three interfaces in the zone:

```
route for 0.0.0.0 0.0.0.0 through outside1 to 10.1.1.2
route for 0.0.0.0 0.0.0.0 through outside2 to 10.2.1.2
route for 0.0.0.0 0.0.0.0 through outside3 to 10.3.1.2
```

Guidelines and Limitations for ECMP

Firewall Mode Guidelines

ECMP zones are supported on routed firewall mode only.

Interface Guidelines

dVTI and Loopback interfaces are not supported.

Additional Guidelines

- A device can have a maximum of 256 ECMP zones.
- You can associate only 8 interfaces per ECMP zone.
- An interface can be a member of only one ECMP zone.
- You cannot remove an interface that is associated with an equal cost static route from the ECMP zone.
- You cannot delete an ECMP zone if its interface has equal cost static routes associated with it.
- Only routed interfaces can be associated with an ECMP zone.
- The following interfaces cannot be associated with an ECMP zone:
  - BVI interface.
  - Member interfaces in an EtherChannel.
  - Failover or state link interface.
  - Management-only or management-access interfaces.
  - Cluster control link interface.
  - VNIs.
  - VLAN interfaces.
  - Interfaces in a remote access VPN configuration with SSL enabled.
- DHCP Relay is not supported on interfaces in an ECMP zone.
- Dual ISP/WAN Firewall Threat Defense Deployment—Create a single ECMP zone for the primary and secondary data interfaces. This configuration enables creation of static routes for both the interfaces with same metric values.
- The Firewall Threat Defense does not support ECMP with NAT in IPsec sessions—a standard IPsec virtual private network (VPN) tunnel does not work with NAT points in the delivery path of IPsec packets.

Manage ECMP Page

When you click **ECMP** on the Routing pane, the ECMP page appears corresponding to the virtual router. This page displays the existing ECMP zones with the associated interfaces of the virtual router. In this page, you can Add an ECMP zone to the virtual router. You can also **Edit** (✎) and **Delete** (🗑) ECMP.

You can perform the following:

- Create an ECMP Zone, on page 1225
- Configure an Equal Cost Static Route, on page 1226
- Modify an ECMP Zone, on page 1227
- Remove an ECMP Zone, on page 1227

Create an ECMP Zone

ECMP zones are created per virtual router. Thus, only the interfaces of the virtual router where the ECMP is being created can be associated with the ECMP.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 From the virtual router drop-down, select the virtual router in which you want to create the ECMP zone.

You can create ECMP zones in global virtual router and user-defined virtual routers. For information on creating virtual routers, see Create a Virtual Router, on page 1178.

Step 4 Click **ECMP**.

Step 5 Click **Add**.

Step 6 In the **Add ECMP** box, enter a name for the ECMP zone.

Note
The ECMP name must be unique for the routed device.

Step 7 To associate interfaces, select the interface under the **Available Interfaces** box, and then click **Add**.

Remember the following:

- Only interfaces belonging to the virtual router are available for assigning.
- Only interfaces with a logical name are listed under the **Available Interfaces** box. You can edit the interface and provide a logical name in **Interfaces**. Remember to save the changes for the settings to take effect.

Step 8 Click **OK**.

The ECMP page now displays the newly created ECMP.

Step 9 Click **Save** and **Deploy** the configuration.

You can associate the ECMP zone interfaces with equal cost static route by defining them with same destination and metric value, but with different gateway.

What to do next

- Configure an Equal Cost Static Route, on page 1226
- Modify an ECMP Zone, on page 1227
- Remove an ECMP Zone, on page 1227

Configure an Equal Cost Static Route

Smart License	Classic License	Supported Devices	Supported Domains	Access
Any	N/A	Firewall Threat Defense and Firewall Threat Defense Virtual	Any	Admin/Network Admin/Security Approver

You can assign interfaces of a virtual router, both global and user-defined, to an ECMP zone for the device.

Before you begin

- To configure an equal cost static route for an interface, ensure to associate it with an ECMP zone. See Create an ECMP Zone, on page 1225.
- All routing configuration settings of a non-VRF capable device are also available for a global virtual router.
- You cannot define a static route for interfaces with same destination and metric without associating the interfaces with an ECMP zone.

Procedure

Step 1 From the **Devices** > **Device Management** page, edit the Firewall Threat Defense device. Click the **Routing** tab.

Step 2 From the drop-down list, select the virtual router whose interfaces are associated with an ECMP zone.

Step 3 To configure the equal cost static route for the interfaces, click **Static Route**.

Step 4 Either click **Add Route** to add a new route, or click **Edit** (✎) for an existing route.

Step 5 From the **Interface** drop-down, select the interface belonging to the virtual router and an ECMP zone.

Step 6 Select the destination network from the **Available Networks** box and click **Add**.

Step 7 Enter a gateway for the network.

Step 8 Enter a metric value. It can be a number that ranges between 1 and 254.

Step 9 To save the settings, click **Save**.

Step 10 To configure equal cost static routing, repeat the steps to configure the static route for another interface in the same ECMP zone with the same destination network and metric value. Remember to provide a different gateway.

What to do next

- Modify an ECMP Zone, on page 1227
- Remove an ECMP Zone, on page 1227

Modify an ECMP Zone

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **ECMP**.

ECMP zones with its associated interfaces are displayed in the **ECMP** page.

Step 4 To modify an ECMP, click **Edit** (✎) against the desired ECMP. In the **Edit ECMP** box, you can do the following:

- ECMP Name—Ensure that the changed name is unique for the device.
- Interfaces—You can add or remove interfaces. You cannot include an interface that is already associated with another ECMP. In addition, you cannot remove the interface that is associated with an equal cost static route.

Step 5 Click **OK**.

Step 6 To save the changes, click **Save**.

What to do next

- Configure an Equal Cost Static Route, on page 1226
- Remove an ECMP Zone, on page 1227

Remove an ECMP Zone

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **ECMP**.

ECMP zones with its associated interfaces are displayed in the **ECMP** page.

Step 4 To remove an ECMP zone, click **Delete** (🗑) against the ECMP zone.

You cannot delete the ECMP zone if any of its interfaces are associated with an equal cost static route.

Step 5 Click **Delete** in the confirmation message.

Step 6 To save the changes, click **Save**.

Configuration Example for ECMP

This example demonstrates how to use Firewall Management Center to configure ECMP zones on Firewall Threat Defense such that the traffic flowing through the device is handled efficiently. With ECMP configured, Firewall Threat Defense maintains the routing table per zone basis, and hence it makes it possible to re-route the packets in the best possible routes. Thus, ECMP supports asymmetric routing, load balancing, and handles lost traffic seamlessly. In this example, R4 records the two paths to reach the external file server.

Figure 416: Configuration Example for ECMP

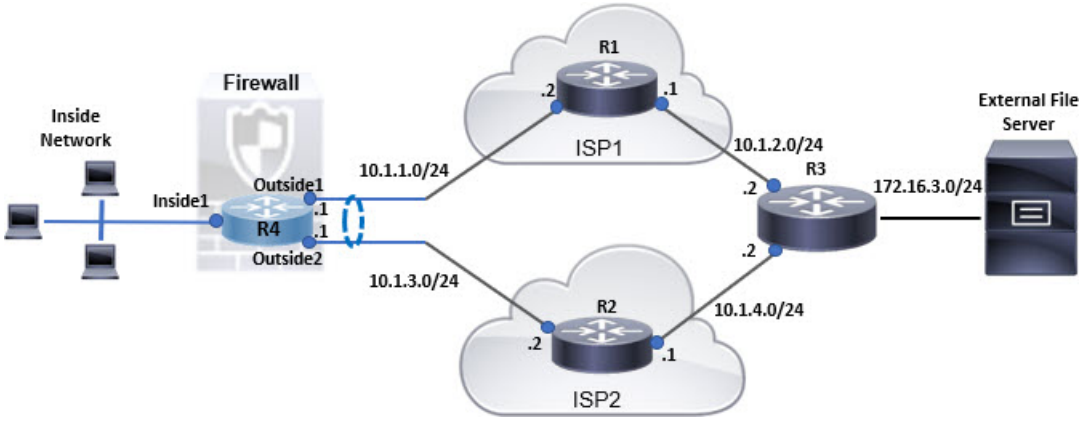

Procedure

Step 1 Create virtual router—*R4* with *Inside1*, *Outside1*, and *Outside2* interfaces:

Figure 417: Configuring R4 Virtual Router

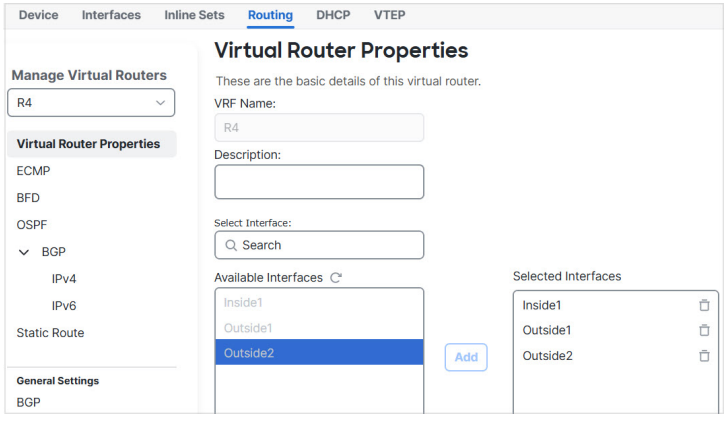

Step 2 Create ECMP zones:

a) In the **Routing** tab, choose R4 user defined virtual router, and then click **ECMP**.

b) Click **Add**.
c) Enter the ECMP name and from the **Available Interfaces** list, choose *Outside1* and *Outside2*:

Figure 418: Creating ECMP Zone

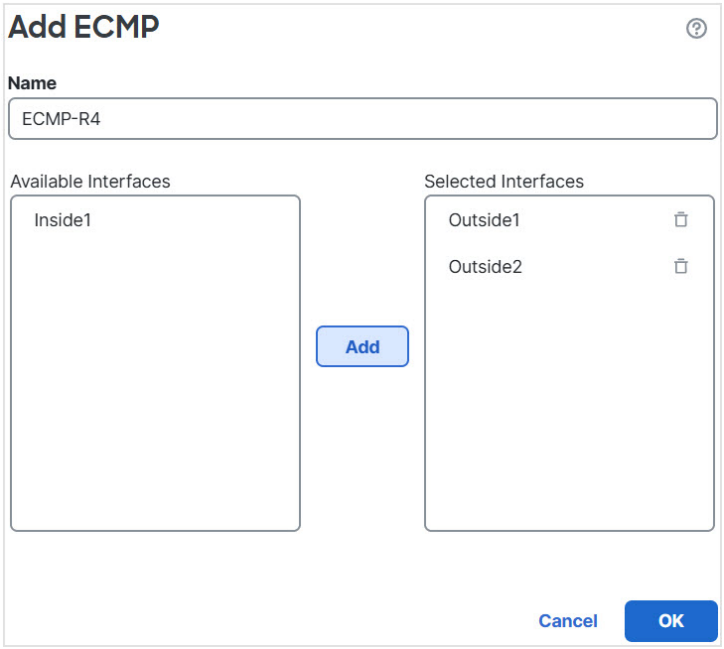

d) Click **Ok**, and then **Save**.

Step 3 Create static routes for the zone interfaces:

a) In the **Routing** tab, click **Static Route**.
b) From the **Interface** drop-down list, select Outside1.
c) Under **Available Network**, choose any-ipv4 and click **Add**.
d) Specify the next-hop address in the **Gateway** field, 10.1.1.2:

Figure 419: Configuring Static Route for Outside1

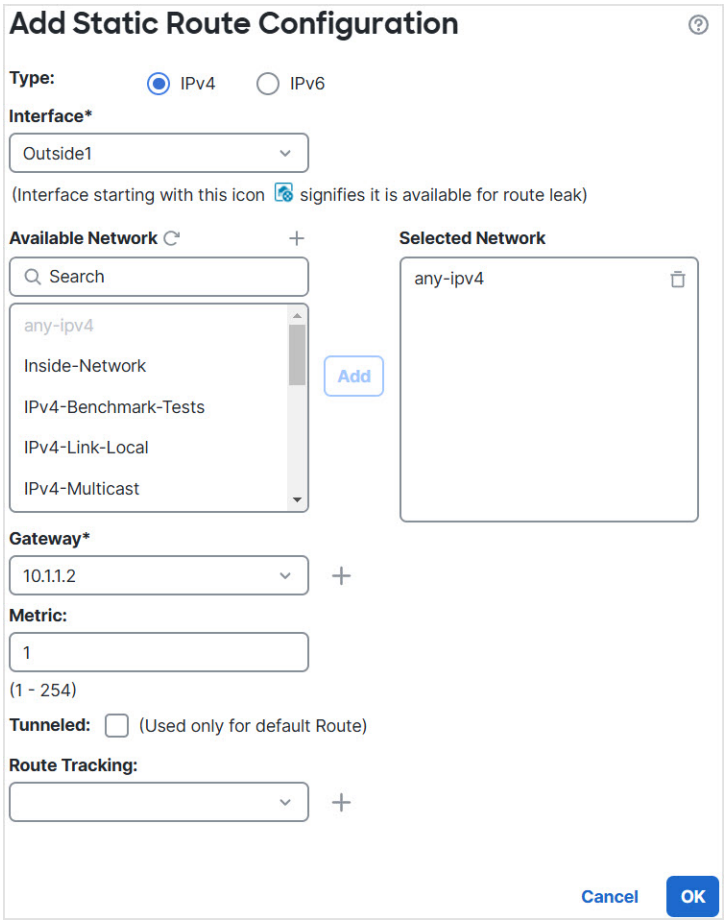

e) Configure the static route for Outside2, repeating from Step 3b to Step 3d.

Ensure to specify same metric, but different gateways for the static routes:

Figure 420: Configured Static Routes of ECMP Zone Interfaces

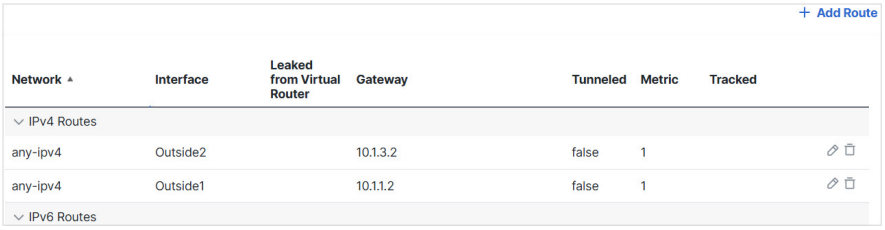

Step 4 **Save** and **Deploy**.

The network packets to reach its destination, R3, follows R4>R1>R3 or R4>R2>R3, based on the ECMP algorithm. If R1>R3 route is lost, the traffic flows through R2 without any packet drops. Similarly, the response from R3 can be received by *Outside2* though the packet was sent from *Outside1*. In addition, when the network traffic is heavy, R4 distributes them between the two routes and thus balances the load.

History for ECMP in Secure Firewall Threat Defense

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
ECMP support as Routing Policy	7.1	Any	Secure Firewall Threat Defense was supporting ECMP routing through FlexConfig policies. From this release, you can group interfaces in to traffic zones and configure ECMP routing in Secure Firewall Management Center. New/modified screens: **Devices** > **Device Management** > **Routing** > **ECMP**

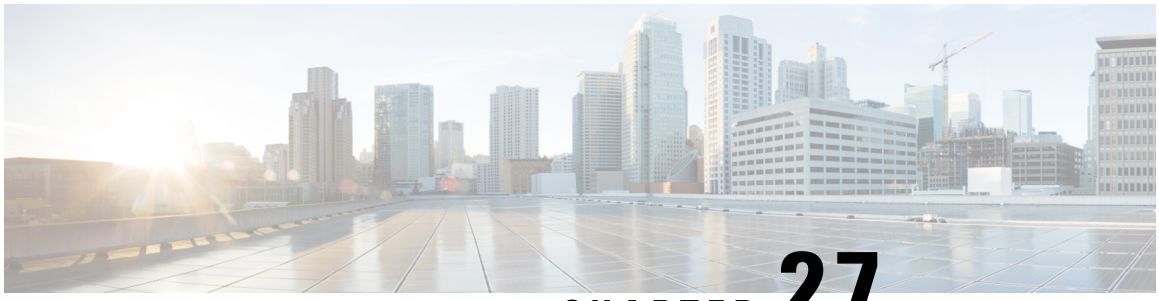

CHAPTER 27

Bidirectional Forwarding Detection Routing

This chapter describes how to configure Firewall Threat Defense to use the Bidirectional Forwarding Detection (BFD) routing protocol.

- About BFD Routing, on page 1233
- Guidelines for BFD Routing, on page 1233
- Configure BFD, on page 1235
- History for BFD Routing, on page 1237

About BFD Routing

BFD is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. BFD operates in a unicast, point-to-point mode on top of any data protocol being forwarded between two systems. Packets are carried in the payload of the encapsulating protocol appropriate for the media and the network.

BFD provides a consistent failure detection method for network administrators in addition to fast forwarding path failure detection. Because the network administrator can use BFD to detect forwarding path failures at a uniform rate, rather than the variable rates for different routing protocol hello mechanisms, network profiling and planning are easier and reconvergence time is consistent and predictable.

Guidelines for BFD Routing

Context Mode Guidelines

BFD is supported on all Firewall Threat Defense platforms. It is supported in multi-instance mode.

Firewall Mode Guidelines

Supported in routed firewall mode and not in transparent mode.

Failover and Cluster Guidelines

- BFD is not supported on failover interfaces.
- In clustering, BFD is supported only on the control node.

Routing and Protocol Guidelines

- OSPFv2, OSPFv3, IS-IS, BGP IPv4, and BGP IPv6 protocol are supported.

Note For optimal routing, do not configure BFD when BGP graceful restart for NSF is configured on the device.

For BFD support on IS-IS, using FlexConfig CLI configure BFD on IS-IS interfaces (Physical Interfaces, Sub-Interfaces, Port-Channels only):

```
For IPv6
###Flex-config Appended CLI###

router isis
  net 11.1111.0000.0000.0001.00
exit
interface GigabitEthernet x/x
  ipv6 router isis
  isis ipv6 bfd
exit

For IPv4
###Flex-config Appended CLI###

router isis
  net 11.1111.0000.0000.0001.00
exit
interface GigabitEthernet x/x
  isis
  isis bfd
exit
```

EIGRP protocol is not supported.

- BFD for static routes is not supported. You can configure BFD on interfaces that belong only to virtual routers.
- Only named interfaces are supported.
- BFD on BVI, VTI, and loopback interfaces are not supported.

Single-hop Guidelines

- Echo mode is disabled by default. You can enable echo mode on single-hop only.
- Echo mode is not supported for IPv6.
- Use only a single-hop template to configure a single-hop policy.
- Authentication of the single-hop template is optional.
- You cannot configure multiple BFDs on the same interface.

Multi-hop Guidelines

- Do not configure the source IP address also as the destination IP address.
- Source and destination address should have same IP type—IPV4 or IPV6.

- Only network objects of host or network type are allowed.
- Use only a multi-hop template to configure a multi-hop policy.
- Authentication is mandatory for the multi-hop template.

Upgrade Guidelines

When you upgrade to version 7.3 and when the previous version has any FlexConfig BFD policies, the management center displays a warning message during deployment. However, it does not stop the deployment process. After post-upgrade deployment, to manage the BFD policies from the UI (**Device (Edit)** > **Routing** > **BFD**), you must configure BFD policies in the **Device (Edit)** > **Routing** > **BFD** page and remove the configuration from the FlexConfig policy for the device.

Configure BFD

This section describes how to enable and configure the BFD routing policy on your system.

Procedure

Step 1 Create BFD Template, on page 1377.

Step 2 Configure BFD Policies, on page 1235.

Step 3 Configure BFD support in the BGP neighbor settings; see, Step 13, on page 1286

Configure BFD Policies

You can bind a BFD template to an interface belonging to a virtual router, or to a source and destination address pair.

Before you begin

- BFD policy is configurable only on interfaces that belong to virtual routers. See Configure interfaces to a virtual router.

Procedure

Step 1 From the **Devices** > **Device Management** page, edit the virtual-router supported device. Navigate to **Routing**.

Step 2 From the drop-down list, select the desired virtual router, and then click **BFD**.

Step 3 To configure a BFD on the interface, click the **Single-Hop** tab or **Multi-Hop** tab.

Note
For a single-hop policy, the BFD template is configured on an interface; for a multi-hop policy, the BFD template is configured on a source and destination address pair.

Step 4 Click **Add**. To modify the configured BFD policy, click **Edit** ( ).

Note
When you edit the interface mapping with BFD template to replace it with a new BFD template, the Firewall Management Center uses a **no** command to remove the template mapping from interface and applies the new template to the interface which causes a BFD flap which may also lead to an OSPFv2, OSPFv3, or BGP flap. However, if the BFD intervals are higher, the BFD flap might not occur. Alternatively, to avoid the flapping, you can delete the existing BFD template mapping; deploy the interface, and then add the new BFD template to the interface and deploy the configuration.

- See Configure Single-Hop BFD Policies, on page 1236.
- See Configure Multi-Hop BFD Policies, on page 1236.

Configure Single-Hop BFD Policies

You can configure a single-hop BFD policy only on an interface that is belonging to a virtual router.

Before you begin

- Create single-hop BFD templates. You cannot configure single-hop BFD policy on interfaces using a multi-hop template.

Procedure

Step 1 In the **Single-Hop** tab, click **Add** or **Edit**.

Step 2 In the **Add BFD Single-Hop** dialog box, configure the following:

a) In the **Interface** drop-down list, interfaces belonging to virtual routers are listed. Select the interface you want to configure with the BFD policy.

b) In the **Template Name** drop-down list, single-hop templates are listed. Select the template that you want to apply.

If you have not created a single-hop template, use **Add** ( ) and create a single-hop BFD template.

Step 3 Click **OK** and **Save** the configuration.

Configure Multi-Hop BFD Policies

You can configure multi-hop BFD policy on a source and destination address pair.

Before you begin

- Create multi-hop BFD templates. You cannot configure multi-hop BFD policy using a single-hop template.

Procedure

Step 1 In the **Add BFD Multi-Hop** dialog box, configure the following:

a) Click the BFD source address type—**IPv4** or **IPv6** radio button.

b) In the **Source Address** drop-down list, network objects are listed. Select the source address that you want to configure for the BFD policy. You cannot choose *any-ipv4* or *any-ipv6*.

If you have not created the required network object, use **Add** (+) and create a host/network object.

Note
The created network object's IP type should match with the selected source IP type.

c) In the **Destination Address** drop-down list, network objects are listed. Select the destination address that you want to configure for the BFD. You cannot choose *any-ipv4* or *any-ipv6*.

If you have not created the required network object, use **Add** (+) and create a host/network object.

Note
The created network object's IP type should match with the selected source IP type.

Attention
Do not select the network object that has the same IP address as that of the source address.

d) In the **Template Name** drop-down list, multi-hop templates are listed. Select the template that you want to apply on the BFD policy.

If you have not created a multi-hop template, use **Add** (+) and create a multi-hop BFD template.

Step 2 Click **OK** and **Save** the configuration.

The multi-hop map (table view) is displayed on the **Multi-Hop** tab page.

History for BFD Routing

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
BFD support for IS-IS	7.4	7.4	You can configure BFD on IS-IS interfaces using FlexConfig CLI.
BFD support for OSPF	7.4	7.4	You can enable BFD on OSPFv2 and OSPFv3 interfaces. New/modified screens: • **Configuration** > **Device Setup** > **Routing** > **OSPFv2** • **Configuration** > **Device Setup** > **Routing** > **OSPFv3**

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
BFD configuration	7.4	7.4	In the previous releases, BFD was configurable on threat defense only through FlexConfig. FlexConfig no longer supports BFD configuration. You can now configure BFD policies for threat defense in the management center UI. In threat defense, BFD is supported only on the BGP protocol. New/modified screens: **Devices** > **Device Management** > **Routing** > **BFD**.

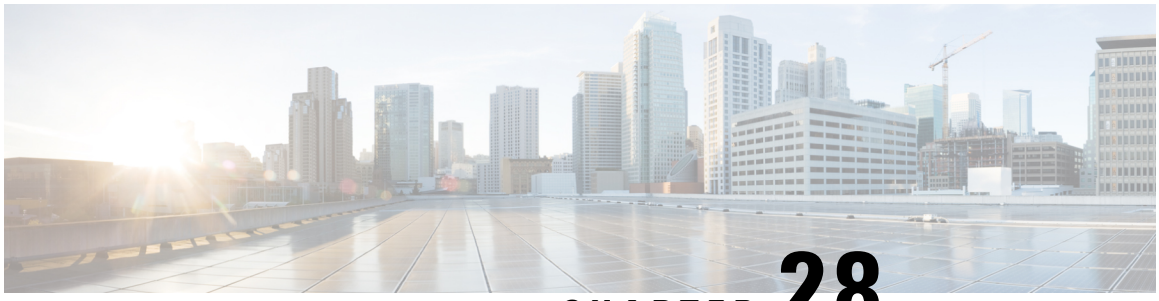

CHAPTER 28

OSPF

This chapter describes how to configure the Firewall Threat Defense to route data, perform authentication, and redistribute routing information using the Open Shortest Path First (OSPF) routing protocol.

- OSPF, on page 1239
- Requirements and Prerequisites for OSPF, on page 1242
- Guidelines for OSPF, on page 1242
- Configure OSPFv2, on page 1245
- Configure OSPFv3, on page 1257
- History for OSPF, on page 1266

OSPF

This chapter describes how to configure the Firewall Threat Defense to route data, perform authentication, and redistribute routing information using the Open Shortest Path First (OSPF) routing protocol.

About OSPF

OSPF is an interior gateway routing protocol that uses link states rather than distance vectors for path selection. OSPF propagates link-state advertisements rather than routing table updates. Because only LSAs are exchanged instead of the entire routing tables, OSPF networks converge more quickly than RIP networks.

OSPF uses a link-state algorithm to build and calculate the shortest path to all known destinations. Each router in an OSPF area contains an identical link-state database, which is a list of each of the router usable interfaces and reachable neighbors.

The advantages of OSPF over RIP include the following:

- OSPF link-state database updates are sent less frequently than RIP updates, and the link-state database is updated instantly, rather than gradually, as stale information is timed out.
- Routing decisions are based on cost, which is an indication of the overhead required to send packets across a certain interface. The Firewall Threat Defense device calculates the cost of an interface based on link bandwidth rather than the number of hops to the destination. The cost can be configured to specify preferred paths.

The disadvantage of shortest path first algorithms is that they require a lot of CPU cycles and memory.

The Firewall Threat Defense device can run two processes of OSPF protocol simultaneously on different sets of interfaces. You might want to run two processes if you have interfaces that use the same IP addresses (NAT allows these interfaces to coexist, but OSPF does not allow overlapping addresses). Or you might want to run one process on the inside and another on the outside, and redistribute a subset of routes between the two processes. Similarly, you might need to segregate private addresses from public addresses.

You can redistribute routes into an OSPF routing process from another OSPF routing process, a RIP routing process, or from static and connected routes configured on OSPF-enabled interfaces.

The Firewall Threat Defense device supports the following OSPF features:

- Intra-area, inter-area, and external (Type I and Type II) routes.
- Virtual links.
- LSA flooding.
- Authentication to OSPF packets (both password and MD5 authentication).
- Configuring the Firewall Threat Defense device as a designated router or a designated backup router. The Firewall Threat Defense device also can be set up as an ABR.
- Stub areas and not-so-stubby areas.
- Area boundary router Type 3 LSA filtering.

OSPF supports MD5 and clear text neighbor authentication. Authentication should be used with all routing protocols when possible because route redistribution between OSPF and other protocols (such as RIP) can potentially be used by attackers to subvert routing information.

If NAT is used, if OSPF is operating on public and private areas, and if address filtering is required, then you need to run two OSPF processes—one process for the public areas and one for the private areas.

A router that has interfaces in multiple areas is called an Area Border Router (ABR). A router that acts as a gateway to redistribute traffic between routers using OSPF and routers using other routing protocols is called an Autonomous System Boundary Router (ASBR).

An ABR uses LSAs to send information about available routes to other OSPF routers. Using ABR Type 3 LSA filtering, you can have separate private and public areas with the ASA acting as an ABR. Type 3 LSAs (inter-area routes) can be filtered from one area to other, which allows you to use NAT and OSPF together without advertising private networks.

Note Only Type 3 LSAs can be filtered. If you configure the Firewall Threat Defense device as an ASBR in a private network, it will send Type 5 LSAs describing private networks, which will get flooded to the entire AS, including public areas.

If NAT is employed but OSPF is only running in public areas, then routes to public networks can be redistributed inside the private network, either as default or Type 5 AS external LSAs. However, you need to configure static routes for the private networks protected by the Firewall Threat Defense device. Also, you should not mix public and private networks on the same Firewall Threat Defense device interface.

You can have two OSPF routing processes, one RIP routing process, and one EIGRP routing process running on the Firewall Threat Defense device at the same time.

OSPF Support for Fast Hello Packets

The OSPF Support for Fast Hello Packets feature provides a way to configure the sending of hello packets in intervals less than one second. Such a configuration would result in faster convergence in an Open Shortest Path First (OSPF) network.

Prerequisites for OSPF Support for Fast Hello Packets

OSPF must be configured in the network already or configured at the same time as the OSPF Support for Fast Hello Packets feature.

OSPF Hello Interval and Dead Interval

OSPF hello packets are packets that an OSPF process sends to its OSPF neighbors to maintain connectivity with those neighbors. The hello packets are sent at a configurable interval (in seconds). The defaults are 10 seconds for an Ethernet link and 30 seconds for a non broadcast link. Hello packets include a list of all neighbors for which a hello packet has been received within the dead interval. The dead interval is also a configurable interval (in seconds), and defaults to four times the value of the hello interval. The value of all hello intervals must be the same within a network. Likewise, the value of all dead intervals must be the same within a network.

These two intervals work together to maintain connectivity by indicating that the link is operational. If a router does not receive a hello packet from a neighbor within the dead interval, it will declare that neighbor to be down.

OSPF Fast Hello Packets

OSPF fast hello packets refer to hello packets being sent at intervals of less than 1 second. To understand fast hello packets, you should already understand the relationship between OSPF hello packets and the dead interval. See OSPF Hello Interval and Dead Interval, on page 1241.

OSPF fast hello packets are achieved by using the ospf dead-interval command. The dead interval is set to 1 second, and the hello-multiplier value is set to the number of hello packets you want sent during that 1 second, thus providing subsecond or "fast" hello packets.

When fast hello packets are configured on the interface, the hello interval advertised in the hello packets that are sent out this interface is set to 0. The hello interval in the hello packets received over this interface is ignored.

The dead interval must be consistent on a segment, whether it is set to 1 second (for fast hello packets) or set to any other value. The hello multiplier need not be the same for the entire segment as long as at least one hello packet is sent within the dead interval.

Benefits of OSPF Fast Hello Packets

The benefit of the OSPF Fast Hello Packets feature is that your OSPF network will experience faster convergence time than it would without fast hello packets. This feature allows you to detect lost neighbors within 1 second. It is especially useful in LAN segments, where neighbor loss might not be detected by the Open System Interconnection (OSI) physical layer and data-link layer.

Implementation Differences Between OSPFv2 and OSPFv3

OSPFv3 is not backward compatible with OSPFv2. To use OSPF to route both IPv4 and IPv6 traffic, you must run both OSPFv2 and OSPFv3 at the same time. They coexist with each other, but do not interact with each other.

The additional features that OSPFv3 provides include the following:

- Protocol processing per link.
- Removal of addressing semantics.
- Addition of flooding scope.
- Support for multiple instances per link.
- Use of the IPv6 link-local address for neighbor discovery and other features.
- LSAs expressed as prefix and prefix length.
- Addition of two LSA types.
- Handling of unknown LSA types.
- Authentication support using the IPsec ESP standard for OSPFv3 routing protocol traffic, as specified by RFC-4552.

Requirements and Prerequisites for OSPF

Model Support

Threat Defense

Firewall Threat Defense Virtual

Supported Domains

Any

User Roles

Admin

Network Admin

Guidelines for OSPF

Firewall Mode Guidelines

OSPF supports routed firewall mode only. OSPF does not support transparent firewall mode.

High Availability Guidelines

OSPFv2 and OSPFv3 support Stateful High Availability.

IPv6 Guidelines

- OSPFv2 does not support IPv6.
- OSPFv3 supports IPv6.
- OSPFv3 uses IPv6 for authentication.
- The Firewall Threat Defense device installs OSPFv3 routes into the IPv6 RIB, provided it is the best route.

OSPFv3 Hello Packets and GRE

Typically, OSPF traffic does not pass through GRE tunnel. When OSPFv3 on IPv6 is encapsulated inside GRE, the IPv6 header validation for security check such as Multicast Destination fails. The packet is dropped due to the implicit security check validation, as this packet has destination IPv6 multicast.

You may define a pre-filter rule to bypass GRE traffic. However, with pre-filter rule, inner packets would not be interrogated by the inspection engine.

Clustering Guidelines

- OSPFv3 encryption is not supported. An error message appears if you try to configure OSPFv3 encryption in a clustering environment.
- In Spanned interface mode, dynamic routing is not supported on management-only interfaces.
- In Individual interface mode, make sure that you establish the control and data units as either OSPFv2 or OSPFv3 neighbors.
- In Individual interface mode, OSPFv2 adjacencies can only be established between two contexts on a shared interface on the control unit. Configuring static neighbors is supported only on point-to point-links; therefore, only one neighbor statement is allowed on an interface.
- When a control role change occurs in the cluster, the following behavior occurs:
  - In spanned interface mode, the router process is active only on the control unit and is in a suspended state on the data units. Each cluster unit has the same router ID because the configuration has been synchronized from the control unit. As a result, a neighboring router does not notice any change in the router ID of the cluster during a role change.
  - In individual interface mode, the router process is active on all the individual cluster units. Each cluster unit chooses its own distinct router ID from the configured cluster pool. A control role change in the cluster does not change the routing topology in any way.

Multiprotocol Label Switching (MPLS) and OSPF Guidelines

When a MPLS-configured router sends Link State (LS) update packets containing opaque Type-10 link-state advertisements (LSAs) that include an MPLS header, authentication fails and the appliance silently drops the update packets, rather than acknowledging them. Eventually the peer router will terminate the neighbor relationship because it has not received any acknowledgments.

Make sure that non-stop forwarding (NSF) is disabled on the appliance to ensure that the neighbor relationship remains stable:

- Navigate to the **Non Stop Forwarding** page in Firewall Management Center(**Devices** > **Device Management (select the desired device)** > **Routing** > **OSPF** > **Advanced** > **Non Stop Forwarding**).

  Ensure the **Non Stop Forwarding Capability** boxes are not checked.

Note The Firepower 4100/9300 models may have high latency when using MPLS because they lack load balancing across multiple receiving queues.

Bidirectional and Forwarding Detection (BFD) and OSPF Guidelines

- You can enable BFD on OSPFv2 and OSPFv3 interfaces (Physical Interfaces, Sub-Interfaces, and Port-Channels).
- BFD is not supported on VTI Tunnels, DVTI Tunnels, Loopback, Switchport, VNI, VTEP, and IRB interfaces.

Route Redistribution Guidelines

- Redistribution of route maps with IPv4 prefix list on OSPFv2 is supported. However, redistribution of route maps with IPv6 prefix list on OSPFv3 is not supported. Use an access list in the route map on OSPF for redistribution.
- When OSPF is configured on a device that is a part of EIGRP network or vice versa, ensure that OSPF-router is configured to tag the route (EIGRP does not support route tag yet).

  When redistributing OSPF into EIGRP and EIGRP into OSPF, a routing loop occurs when there is an outage on one of the links, interfaces, or even when the route originator is down. To prevent the redistribution of routes from one domain back into the same domain, a router can tag a route that belongs to a domain while it is redistributing, and those routes can be filtered on the remote router based on the same tag. Because the routes will not be installed into the routing table, they will not be redistributed back into the same domain.

Additional Guidelines

- OSPFv2 and OSPFv3 support multiple instances on an interface.
- OSPFv3 supports encryption through ESP headers in a non-clustered environment.
- OSPFv3 supports Non-Payload Encryption.
- OSPFv2 supports Cisco NSF Graceful Restart and IETF NSF Graceful Restart mechanisms as defined in RFCs 4811, 4812 & 3623 respectively.
- OSPFv3 supports Graceful Restart mechanism as defined in RFC 5187.
- There is a limit to the number of intra area (type 1) routes that can be distributed. For these routes, a single type-1 LSA contains all prefixes. Because the system has a limit of 35 KB for packet size, 3000 routes result in a packet that exceeds the limit. Consider 2900 type 1 routes to be the maximum number supported.

- For a device using virtual routing, you can configure OSPFv2 and OSPFv3 for a global virtual router. However, you can configure only OSPFv2 for a user-defined virtual router.
- To avoid adjacency flaps due to route updates being dropped if the route update is larger than the minimum MTU on the link, ensure that you configure the same MTU on the interfaces on both sides of the link.
- OSPFv3 drops the LS update if the packet size exceeds 8190. As a result, the neighbourship is terminated. Hence, configure the switch with the "ospfv3 mtu-ignore" command to avoid neighbourship termination.

Configure OSPFv2

This section describes the tasks involved in configuring an OSPFv2 routing process. For a device using virtual routing, you can configure OSPFv2 for global as well as for user-defined virtual routers.

Configure OSPF Areas, Ranges, and Virtual Links

You can configure several OSPF area parameters, which include setting authentication, defining stub areas, and assigning specific costs to the default summary route. You can enable up to two OSPF process instances. Each OSPF process has its own associated areas and networks. Authentication provides password-based protection against unauthorized access to an area.

Stub areas are areas into which information on external routes is not sent. Instead, there is a default external route generated by the ABR into the stub area for destinations outside the autonomous system. To take advantage of the OSPF stub area support, default routing must be used in the stub area.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.

Step 4 Click **OSPF**.

Step 5 Check the check box of **Process 1**. You can enable up to two OSPF process instances for each context/virtual router. You must choose an OSPF process to be able to configure the Area parameters.

If the device is using virtual routing, the ID fields display the unique process IDs generated for the chosen virtual router.

Step 6 Choose the **OSPF Role** from the drop-down list, and enter a description for it in the next field. The options are Internal, ABR, ASBR, and ABR & ASBR. See About OSPF, on page 1239 for a description of the OSPF roles.

Step 7 Select **Area** > **Add**.

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 8 Configure the following area options for each OSPF process:

- **OSPF Process**—Choose the process ID. For a device using virtual routing, the drop-down lists the unique process IDs generated for the selected virtual router.

- **Area ID**—Designation of the area for which routes are to be summarized.
- **Area Type**—Choose one of the following:
  - **Normal**—(Default) Standard OSPF area.
  - **Stub**—A stub area does not have any routers or areas beyond it. Stub areas prevent Autonomous System (AS) External LSAs (Type 5 LSAs) from being flooded into the stub area. When you create a stub area, you can prevent summary LSAs (Types 3 and 4) from being flooded into the area by NOT checking the **Summary Stub** check box.
  - **NSSA**—Makes the area a not-so-stubby area (NSSA). NSSAs accept Type 7 LSAs. You can disable route redistribution by NOT checking the **Redistribute** check box and checking the **Default Information Originate** check box. You can prevent summary LSAs from being flooded into the area by NOT checking the **Summary NSSA** check box.
- **Metric Value**—The metric used for generating the default route. The default value is 10. Valid metric values range from 0 through 16777214.
- **Metric Type**—The metric type is the external link type that is associated with the default route that is advertised into the OSPF routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route.
- **Available Network**—Choose one of the available networks and click **Add**, or click **Add** (+) to add a new network object. See Network, on page 1401 for the procedure for adding networks.
- **Authentication**—Choose the OSPF authentication:
  - **None**—(Default) Disables OSPF area authentication.
  - **Password**—Provides a clear text password for area authentication, which is not recommended where security is a concern.
  - **MD5**—Allows MD5 authentication.
- **Default Cost**—The default cost for the OSPF area, which is used to determine the shortest paths to the destination. Valid values range from 0 through 65535. The default value is 1.

Step 9 Click **OK** to save the area configuration.

Step 10 Select **Range** > **Add**.

- Choose one of the available networks and whether to advertise, or,
- Click **Add** (+) to add a new network object. See Network, on page 1401 for the procedure for adding networks.

Step 11 Click **OK** to save the range configuration.

Step 12 Select **Virtual Link**, click **Add** (+), and configure the following options for each OSPF process:

- **Peer Router**—Choose the IP address of the peer router. To add a new peer router, click **Add** (+). See Network, on page 1401 for the procedure for adding networks.
- **Hello Interval**—The time in seconds between the hello packets sent on an interface. The hello interval is an unsigned integer that is to be advertised in the hello packets. The value must be the same for all

routers and access servers on a specific network. Valid values range from 1 through 65535. The default is 10.

The smaller the hello interval, the faster topological changes are detected, but the more traffic is sent on the interface.

- **Transmit Delay**—The estimated time in seconds that is required to send an LSA packet on the interface. The integer value must be greater than zero. Valid values range from 1 through 8192. The default is 1.

  LSAs in the update packet have their own ages incremented by this amount before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links.

- **Retransmit Interval**—The time in seconds between LSA retransmissions for adjacencies that belong to the interface. The retransmit interval is the expected round-trip delay between any two routers on the attached network. The value must be greater than the expected round-trip delay, and can range from 1 through 65535. The default is 5.

  When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links.

- **Dead Interval**—The time in seconds that hello packets are not seen before a neighbor indicates that the router is down. The dead interval is an unsigned integer. The default is four times the hello interval, or 40 seconds. The value must be the same for all routers and access servers that are attached to a common network. Valid values range from 1 through 65535.

- **Authentication**—Choose the OSPF virtual link authentication from the following:

  - **None**—(Default) Disables virtual link area authentication.
  - **Area Authentication**—Enables area authentication using MD5. Click **Add**, and enter the key ID, key, confirm the key, and then click **OK**.
  - **Password**—Provides a clear text password for virtual link authentication, which is not recommended where security is a concern.
  - **MD5**—Allows MD5 authentication. Click **Add**, and enter the key ID, key, confirm the key, and then click **OK**.

    Note
    Ensure to enter only numbers as the MD5 key ID.

  - **Key Chain**—Allows key chain authentication. Click **Add**, and create the key chain, and then click **Save**. For detailed procedure, see Creating Key Chain Objects, on page 1399. Use the same authentication type (MD5 or Key Chain) and key ID for the peers to establish a successful adjacency.

Step 13 Click **OK** to save the virtual link configuration.

Step 14 Click **Save** on the Routing page to save your changes.

What to do next

Continue with Configure OSPF Redistribution.

Configure OSPF Redistribution

The Firewall Threat Defense device can control the redistribution of routes between the OSPF routing processes. The rules for redistributing routes from one routing process into an OSPF routing process are displayed. You can redistribute routes discovered by EIGRP, RIP and BGP into the OSPF routing process. You can also redistribute static and connected routes into the OSPF routing process.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.

Step 4 Click **OSPF**.

Step 5 From **OSPF Role** drop-down, choose role .

Step 6 Click **Redistribution** > **Add**.

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 7 Configure the following redistribution options for each OSPF process:

- **OSPF Process**—Choose the process ID. For a device using virtual routing, this drop-down list displays the unique process IDs generated for the selected virtual router.
- **Route Type**—Choose one of the following types:
  - **Static**—Redistributes static routes to the OSPF routing process.
  - **Connected**—Redistributes connected routes (routes established automatically by virtue of having the IP address enabled on the interface) to the OSPF routing process. Connected routes are redistributed as external to the device. You can select whether to use subnets under the Optional list.
  - **OSPF**—Redistributes routes from another OSPF routing process, for example, internal, external 1 and 2, NSSA external 1 and 2, or whether to use subnets. You can select these options under the Optional list.
  - **BGP**—Redistribute routes from the BGP routing process. Add the AS number and whether to use subnets.
  - **RIP**—Redistributes routes from the RIP routing process. You can select whether to use subnets under the Optional list.

    Note
    As a user-defined virtual router does not support RIP, you cannot redistribute routes from RIP.

  - **EIGRP**—Redistribute routes from the EIGRP routing process. Add the AS number and whether to use subnets.
- **Metric Value**—Metric value for the routes being distributed. The default value is 10. Valid values range from 0 to 16777214.

When redistributing from one OSPF process to another OSPF process on the same device, the metric will be carried through from one process to the other if no metric value is specified. When redistributing other processes to an OSPF process, the default metric is 20 when no metric value is specified.

- **Metric Type**—The metric type is the external link type that is associated with the default route that is advertised into the OSPF routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route.
- **Tag Value**—Tag specifies the 32-bit decimal value attached to each external route that is not used by OSPF itself, but which may be used to communicate information between ASBRs. If none is specified, then the remote autonomous system number is used for routes from BGP and EGP. For other protocols, zero is used. Valid values are from 0 to 4294967295.
- **RouteMap**—Checks for filtering the importing of routes from the source routing protocol to the current routing protocol. If this parameter is not specified, all routes are redistributed. If this parameter is specified, but no route map tags are listed, no routes are imported. Or you can add a new route map by clicking **Add** (+). See Configure Route Map Entry to add a new route map.

Step 8 Click **OK** to save the redistribution configuration.

Step 9 Click **Save** on the Routing page to save your changes.

What to do next

Continue with Configure OSPF Inter-Area Filtering, on page 1249.

Configure OSPF Inter-Area Filtering

ABR type 3 LSA filtering extends the capability of an ABR that is running OSPF to filter type 3 LSAs between different OSPF areas. Once a prefix list is configured, only the specified prefixes are sent from one OSPF area to another OSPF area. All other prefixes are restricted to their OSPF area. You can apply this type of area filtering to traffic going into or coming out of an OSPF area, or to both the incoming and outgoing traffic for that area.

When multiple entries of a prefix list match a given prefix, the entry with the lowest sequence number is used. For efficiency, you may want to put the most common matches or denials near the top of the list by manually assigning them a lower sequence number. By default, sequence numbers are automatically generated in increments of 5, beginning with 5.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.

Step 4 Click **OSPF**.

Step 5 Select **InterArea** > **Add**.

You can click **Edit** (✎), or use the right-click menu to cut, copy, past, insert, and delete inter-areas.

Step 6 Configure the following inter-area filtering options for each OSPF process:

- **OSPF Process**—For a device using virtual routing, the drop-down lists the unique process IDs generated for the selected virtual router.
- **Area ID**—The area for which routes are to be summarized.
- **PrefixList**—The name of the prefix. To add a new prefix list object, see Step 5.
- **Traffic Direction**—Inbound or outbound. Choose Inbound to filter LSAs coming into an OSPF area, or Outbound to filter LSAs coming out of an OSPF area. If you are editing an existing filter entry, you cannot modify this setting.

Step 7 Click **Add** (+), and enter a name for the new prefix list, and whether to allow overrides.

You must configure a prefix list before you can configure a prefix rule.

Step 8 Click **Add** to configure prefix rules, and configure the following parameters:

- **Action**—Select **Block** or **Allow** for the redistribution access.
- **Sequence No**—The routing sequence number. By default, sequence numbers are automatically generated in increments of 5, beginning with 5.
- **IP Address**—Specify the prefix number in the format of IP address/mask length.
- **Min Prefix Length**—(Optional) The minimum prefix length.
- **Max Prefix Length**—(Optional) The maximum prefix length.

Step 9 Click **OK** to save the inter-area filtering configuration.

Step 10 Click **Save** on the Routing page to save your changes.

What to do next

Continue with Configure OSPF Filter Rules, on page 1250.

Configure OSPF Filter Rules

You can configure ABR Type 3 LSA filters for each OSPF process. ABR Type 3 LSA filters allow only specified prefixes to be sent from one area to another area and restrict all other prefixes. You can apply this type of area filtering out of a specific OSPF area, into a specific OSPF area, or into and out of the same OSPF area at the same time. OSPF ABR Type 3 LSA filtering improves your control of route distribution between OSPF areas.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.

Step 4 Click **OSPF**.

Step 5 Select **Filter Rule** > **Add**.

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete filter rules.

Step 6 Configure the following filter rule options for each OSPF process:

- **OSPF Process**— For a device using virtual routing, the drop-down lists the unique process IDs generated for the selected virtual router.
- **Access List**—The access list for this OSPF process. To add a new standard access list object, click **Add** () and see Configure Standard ACL Objects, on page 1374.
- **Traffic Direction**—Choose In or Out for the traffic direction being filtered. Choose In to filter LSAs coming into an OSPF area, or Out to filter LSAs coming out of an OSPF area. If you are editing an existing filter entry, you cannot modify this setting.
- **Interface**—The interface for this filter rule.

Step 7 Click **OK** to save the filter rule configuration.

Step 8 Click **Save** on the Routing page to save your changes.

What to do next

Continue with Configure OSPF Summary Addresses, on page 1251.

Configure OSPF Summary Addresses

When routes from other protocols are redistributed into OSPF, each route is advertised individually in an external LSA. However, you can configure the Firewall Threat Defense device to advertise a single route for all the redistributed routes that are included for a specified network address and mask. This configuration decreases the size of the OSPF link-state database. Routes that match the specified IP address mask pair can be suppressed. The tag value can be used as a match value for controlling redistribution through route maps.

Routes learned from other routing protocols can be summarized. The metric used to advertise the summary is the smallest metric of all the more specific routes. Summary routes help reduce the size of the routing table.

Using summary routes for OSPF causes an OSPF ASBR to advertise one external route as an aggregate for all redistributed routes that are covered by the address. Only routes from other routing protocols that are being redistributed into OSPF can be summarized.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.

Step 4 Click **OSPF**.

Step 5 Select **Summary Address** > **Add**.

You can click **Edit** (✎) to edit, or use the right-click menu to cut, copy, past, insert, and delete summary addresses.

Step 6 Configure the following summary address options for each OSPF process:

- **OSPF Process**— For a device using virtual routing, the drop-down lists the unique process IDs generated for the selected virtual router.
- **Available Network**—The IP address of the summary address. Select one from the Available networks list and click **Add**, or to add a new network, click **Add** (+). See Network, on page 1401 for the procedure for adding networks.
- **Tag**—A 32-bit decimal value that is attached to each external route. This value is not used by OSPF itself, but may be used to communicate information between ASBRs.
- **Advertise—Advertises** the summary route. Uncheck this check box to suppress routes that fall under the summary address. By default, this check box is checked.

Step 7 Click **OK** to save the summary address configuration.

Step 8 Click **Save** on the Routing page to save your changes.

What to do next

Continue with Configure OSPF Interfaces and Neighbors, on page 1252.

Configure OSPF Interfaces and Neighbors

You can change some interface-specific OSPFv2 parameters, if necessary. You are not required to change any of these parameters, but the following interface parameters must be consistent across all routers in an attached network: the hello interval, the dead interval, and the authentication key. If you configure any of these parameters, be sure that the configurations for all routers on your network have compatible values.

You need to define static OSPFv2 neighbors to advertise OSPFv2 routes over a point-to-point, non-broadcast network. This feature lets you broadcast OSPFv2 advertisements across an existing VPN connection without having to encapsulate the advertisements in a GRE tunnel.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.

Step 4 Click **OSPF**.

Step 5 Select **Interface** > **Add**.

You can click **Edit** (✎), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 6 Configure the following Interface options for each OSPF process:

- **Interface**—The interface you are configuring.

  Note
  If the device is using virtual routing, this drop-down list displays only those interfaces that belong to the router.

- **Default Cost**—The cost of sending a packet through the interface. The default value is 10.
- **Priority**—The designated router for a network. Valid values range from 0 to 255. The default value is 1. Entering 0 for this setting makes the router ineligible to become the designated router or backup designated router.

  When two routers connect to a network, both attempt to become the designated router. The device with the higher router priority becomes the designated router. If there is a tie, the router with the higher router ID becomes the designated router. This setting does not apply to interfaces that are configured as point-to-point interfaces.

- **MTU Ignore—OSPF** checks whether neighbors are using the same MTU on a common interface. This check is performed when neighbors exchange DBD packets. If the receiving MTU in the DBD packet is higher than the IP MTU configured on the incoming interface, OSPF adjacency is not established.
- **Database Filter**—Use this setting to filter the outgoing LSA interface during synchronization and flooding. By default, OSPF floods new LSAs over all interfaces in the same area, except the interface on which the LSA arrives. In a fully meshed topology, this flooding can waste bandwidth and lead to excessive link and CPU usage. Checking this check box prevents OSPF flooding of the LSA on the selected interface.
- **Hello Interval**—Specifies the interval, in seconds, between hello packets sent on an interface. Valid values range 1–8192 seconds. The default value is 10 seconds.

  The smaller the hello interval, the faster topological changes are detected, but more traffic is sent on the interface. This value must be the same for all routers and access servers on a specific interface.

- **Transmit Delay**—Estimated time in seconds to send an LSA packet on the interface. Valid values range 1–65535 seconds. The default is 1 second.

  LSAs in the update packet have their ages increased by the amount specified by this field before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links.

- **Retransmit Interval**—Time in seconds between LSA retransmissions for adjacencies that belong to the interface. The time must be greater than the expected round-trip delay between any two routers on the attached network. Valid values range from 1 to 65535 seconds. The default is 5 seconds.

  When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links.

- **Dead Interval**—Time period in seconds for which hello packets must not be seen before neighbors indicate that the router is down. The value must be the same for all nodes on the network and can range 1–65535.
- **Hello Multiplier**—Specifies the number of Hello packets to be sent per second. Valid values are 3–20.

- **Point-to-Point**—Lets you transmit OSPF routes over VPN tunnels.
- **Authentication**—Choose the OSPF interface authentication from the following:
  - **None**—(Default) Disables interface authentication.
  - **Area Authentication**—Enables interface authentication using MD5. Click **Add**, and enter the key ID, key, confirm the key, and then click **OK**.
  - **Password**—Provides a clear text password for virtual link authentication, which is not recommended where security is a concern.
  - **MD5**—Allows MD5 authentication. Click **Add**, and enter the key ID, key, confirm the key, and then click **OK**.

    Note
    Ensure to enter only numbers as the MD5 key ID.
  - **Key Chain**—Allows key chain authentication. Click **Add**, and create the key chain, and then click **Save**. For detailed procedure, see Creating Key Chain Objects, on page 1399. Use the same authentication type (MD5 or Key Chain) and key ID for the peers to establish a successful adjacency.
- **Enable BFD**—Allows you to enable BFD on this interface.
- **Enter Password**—The password you configure if you choose Password as the type of authentication.
- **Confirm Password**—Confirm the password that you chose.

Step 7 Select **Neighbor** > **Add**.

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 8 Configure the following parameters for each OSPF process:

- **OSPF Process**—Choose 1 or 2.
- **Neighbor**—Choose one of the neighbors in the drop-down list, or click **Add** () to add a new neighbor; enter the name, description, network, whether to allow overrides, and then click **Save**.
- **Interface**—Choose the interface associated with the neighbor.

Step 9 Click **OK** to save the neighbor configuration.

Step 10 Click **Save** on the Routing page to save your changes.

Configure OSPF Advanced Properties

The Advanced Properties allows you to configure options, such as syslog message generation, administrative route distances, an LSA timer, and graceful restarts.

Graceful Restarts

The Firewall Threat Defense device may experience some known failure situations that should not affect packet forwarding across the switching platform. The Non-Stop Forwarding (NSF) capability allows data forwarding to continue along known routes, while the routing protocol information is being restored.

This capability is useful when there is a scheduled hitless software upgrade. You can configure graceful restart on OSPFv2 by using either using NSF Cisco (RFC 4811 and RFC 4812) or NSF IETF (RFC 3623).

Note NSF capability is also useful in HA mode and clustering.

Configuring the NSF graceful-restart feature involves two steps; configuring capabilities and configuring a device as NSF-capable or NSF-aware. A NSF-capable device can indicate its own restart activities to neighbors and a NSF-aware device can help a restarting neighbor.

A device can be configured as NSF-capable or NSF-aware, depending on some conditions:

- A device can be configured as NSF-aware irrespective of the mode in which it is.
- A device has to be in either Failover or Spanned Etherchannel (L2) cluster mode to be configured as NSF-capable.
- For a device to be either NSF-aware or NSF-capable, it should be configured with the capability of handling opaque Link State Advertisements (LSAs)/ Link Local Signaling (LLS) block as required.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring OSPF.

Step 4 Click **OSPF** > **Advanced**.

Step 5 Select **General**, and configure the following:

- **Router ID**—Choose Automatic or IP Address (appears for non-cluster and a cluster in spanned etherchannel mode) or Cluster Pool (appears for a cluster in individual interface mode) for the router ID. If you choose IP address, enter the IP address in the adjacent field. If you choose Cluster Pool, choose the IPv4 cluster pool value in the adjacent drop-down field. For information on creating the cluster pool address, see Address Pools, on page 1374.
- **Ignore LSA MOSPF**—Suppresses syslog messages when the route receives unsupported LSA Type 6 multicast OSPF (MOSPF) packets.
- **RFC 1583 Compatible**—Configures RFC 1583 compatibility as the method used to calculate summary route costs. Routing loops can occur with RFC 1583 compatibility enabled. Disable it to prevent routing loops. All OSPF routers in an OSPF routing domain should have RFC compatibility set identically.
- **Adjacency Changes**—Defines the adjacency changes that cause syslog messages to be sent.

  By default, a syslog message is generated when an OSPF neighbor goes up or down. You can configure the router to send a syslog message when an OSPF neighbor goes down and also a syslog for each state.

  - **Log Adjacency Changes**—Causes the Firewall Threat Defense device to send a syslog message whenever an OSPF neighbor goes up or down. This setting is checked by default.

- **Log Adjacency Change Details**—Causes the Firewall Threat Defense device to send a syslog message whenever any state change occurs, not just when a neighbor goes up or down. This setting is unchecked by default.

- **Administrative Route Distance**—Allows you to modify the settings that were used to configure administrative route distances for **inter-area**, **intra-area**, and **external** IPv6 routes. The administrative route distance is an integer from 1 to 254. The default is 110.
- **LSA Group Pacing**—Specifies the interval in seconds at which LSAs are collected into a group and refreshed, check summed, or aged. Valid values range from 10 to 1800. The default value is 240.
- **Enable Default Information Originate**—Check the **Enable** check box to generate a default external route into an OSPF routing domain and configure the following options:
  - **Always advertise the default route**—Ensures that the default route is always advertised.
  - **Metric Value**—Metric used for generating the default route. Valid metric values range from 0 to 16777214. The default value is 10.
  - **Metric Type**—The external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. Valid values are 1 (Type 1 external route) and 2 (Type 2 external route). The default is Type 2 external route.
  - **RouteMap**—Choose the routing process that generates the default route if the route map is satisfied or click **Add** (+) to add a new one. See Configure Route Map Entry to add a new route map.

Step 6 Click **OK** to save the general configuration.

Step 7 Select **Non Stop Forwarding**, and configure Cisco NSF graceful restart for OSPFv2, for an NSF-capable or NSF-aware device:

Note
There are two graceful restart mechanisms for OSPFv2, Cisco NSF and IETF NSF. Only one of these graceful restart mechanisms can be configured at a time for an OSPF instance. An NSF-aware device can be configured as both Cisco NSF helper and IETF NSF helper but a NSF-capable device can be configured in either Cisco NSF or IETF NSF mode at a time for an OSPF instance.

a) Check the **Enable Cisco Non Stop Forwarding Capability** check box.
b) (Optional) Check the **Cancel NSF restart when non-NSF-aware neighboring networking devices are detected** check box if required.
c) (Optional) Make sure the **Enable Cisco Non Stop Forwarding Helper** mode check box is unchecked to disable the helper mode on an NSF-aware device.

Step 8 Configure IETF NSF Graceful Restart for OSPFv2, for an NSF-capable or NSF-aware device:

a) Check the **Enable IETF Non Stop Forwarding Capability** check box.
b) In the **Length of graceful restart interval (seconds)** field, enter the restart interval in seconds. The default value is 120 seconds. For a restart interval below 30 seconds, graceful restart will be terminated.
c) (Optional) Make sure the **Enable IETF nonstop forwarding (NSF) for helper mode** check box is unchecked to disable the IETF NSF helper mode on an NSF-aware device.
d) **Enable Strict Link State advertisement checking**—When enabled, it indicates that the helper router will terminate the process of restarting the router if it detects that there is a change to a LSA that would be flooded to the restarting router, or if there is a changed LSA on the retransmission list of the restarting router when the graceful restart process is initiated.

e) **Enable IETF Non Stop Forwarding**—Enables non stop forwarding, which allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following a switchover. OSPF uses extensions to the OSPF protocol to recover its state from neighboring OSPF devices. For the recovery to work, the neighbors must support the NSF protocol extensions and be willing to act as "helpers" to the device that is restarting. The neighbors must also continue forwarding data traffic to the device that is restarting while protocol state recovery takes place.

Configure OSPFv3

This section describes the tasks involved in configuring an OSPFv3 routing process. For a device using virtual routing, you can configure OSPFv3 only for its global virtual router and not for its user-defined virtual router.

Configure OSPFv3 Areas, Route Summaries, and Virtual Links

To enable OSPFv3, you need to create an OSPFv3 routing process, create an area for OSPFv3, enable an interface for OSPFv3, and then redistribute the route into the targeted OSPFv3 routing process.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select **Routing** > **OSPFv3**.

Step 3 By default **Enable Process 1** is selected. You can enable up to two OSPF process instances.

Step 4 Chose the OSPFv3 role from the drop-down list, and enter a description for it. The options are Internal, ABR, ASBR, and ABR and ASBR. See About OSPF, on page 1239 for descriptions of the OSPFv3 roles.

Step 5 Select **Area** > **Add**.

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 6 Select **General**, and configure the following options for each OSPF process:

- **Area ID**—The area for which routes are to be summarized.
- **Cost**—The metric or cost for the summary route, which is used during OSPF SPF calculations to determine the shortest paths to the destination. Valid values range from 0 to 16777215.
- **Type**—Specifies Normal, NSSA, or Stub. If you select Normal, there are no other parameters to configure. If you select Stub, you can choose to send summary LSAs in the area. If you select NSSA, you can configure the next three options:
  - **Allow Sending summary LSA into this area**—Allows the sending of summary LSAs into the area.
  - **Imports routes to normal and NSSA area**—Allows redistribution to import routes to normal and not to stubby areas.
  - **Defaults information originate**—Generates a default external route into an OSPFv3 routing domain.

- **Metric**—Metric used for generating the default route. The default value is 10. Valid metric values range from 0 to 16777214.
- **Metric Type**—The metric type is the external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route.

Step 7 Click **OK** to save the general configuration.

Step 8 (Not applicable for Internal OSPFv3 Role) Select **Route Summary** > **Add Route Summary**.

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete route summaries.

Step 9 Configure the following route summary options for each OSPF process:

- **IPv6 Prefix/Length**—The IPv6 prefix. To add a new network object, click **Add** (). See Network, on page 1401 for the procedure for adding networks.
- **Cost**—The metric or cost for the summary route, which is used during OSPF SPF calculations to determine the shortest paths to the destination. Valid values range from 0 to 16777215.
- **Advertise**—Advertises the summary route. Uncheck this check box to suppress routes that fall under the summary address. By default, this check box is checked.

Step 10 Click **OK** to save the route summary configuration.

Step 11 (Not applicable for Internal OSPFv3 Role) Select **Virtual Link**, click **Add Virtual Link**, and configure the following options for each OSPF process:

- **Peer RouterID**—Choose the IP address of the peer router. To add a new network object, click **Add** (). See Network, on page 1401 for the procedure for adding networks.
- **TTL Security**—Enables TTL security check. The value for the hop-count is a number from 1 to 254. The default is 1.

  OSPF sends outgoing packets with an IP header Time to Live (TTL) value of 255 and discards incoming packets that have TTL values less than a configurable threshold. Because each device that forwards an IP packet decrements the TTL, packets received via a direct (one-hop) connection have a value of 255. Packets that cross two hops have a value of 254, and so on. The receive threshold is configured in terms of the maximum number of hops that a packet may have traveled.
- **Dead Interval**—The time in seconds that hello packets are not seen before a neighbor indicates that the router is down. The default is four times the hello interval, or 40 seconds. Valid values range from 1 to 65535.

  The dead interval is an unsigned integer. The value must be the same for all routers and access servers that are attached to a common network.
- **Hello Interval**—The time in seconds between the hello packets sent on an interface. Valid values range from 1 to 65535. The default is 10.

  The hello interval is an unsigned integer that is to be advertised in the hello packets. The value must be the same for all routers and access servers on a specific network. The smaller the hello interval, the faster topological changes are detected, but the more traffic is sent on the interface.
- **Retransmit Interval**—The time in seconds between LSA retransmissions for adjacencies that belong to the interface. The retransmit interval is the expected round-trip delay between any two routers on the

attached network. The value must be greater than the expected round-trip delay, and can range from 1 to 65535. The default is 5.

When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links.

- **Transmit Delay**—The estimated time in seconds that is required to send an LSA packet on the interface. The integer value must be greater than zero. Valid values range from 1 to 8192. The default is 1.

  LSAs in the update packet have their own ages incremented by this amount before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links.

Step 12 Click **OK** to save the virtual link configuration.

Step 13 Click **Save** on the Router page to save your changes.

What to do next

Continue with Configure OSPFv3 Redistribution.

Configure OSPFv3 Redistribution

The Secure Firewall Threat Defense device can control the redistribution of routes between the OSPF routing processes. The rules for redistributing routes from one routing process into an OSPF routing process are displayed. You can redistribute routes discovered by EIGRP, RIP and BGP into the OSPF routing process. You can also redistribute static and connected routes into the OSPF routing process.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select **Routing** > **OSPF**.

Step 3 Select **Redistribution**, and click **Add.**

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 4 Configure the following redistribution options for each OSPF process:

- **Source Protocol**—The source protocol from which routes are being redistributed. The supported protocols are connected, OSPF, Static, EIGRP, and BGP. If you choose OSPF, you must enter the Process ID in the **Process ID** field. If you choose BGP, you must add the AS number in the **AS Number** field.
- **Metric** —Metric value for the routes being distributed. The default value is 10. Valid values range from 0 to 16777214.

  When redistributing from one OSPF process to another OSPF process on the same device, the metric will be carried through from one process to the other if no metric value is specified. When redistributing other processes to an OSPF process, the default metric is 20 when no metric value is specified.

- **Metric Type**—The metric type is the external link type that is associated with the default route that is advertised into the OSPF routing domain. The available options are 1 for a Type 1 external route or 2 for a Type 2 external route.
- **Tag** —Tag specifies the 32-bit decimal value attached to each external route that is not used by OSPF itself, but which may be used to communicate information between ASBRs. If none is specified, then the remote autonomous system number is used for routes from BGP and EGP. For other protocols, zero is used. Valid values are from 0 to 4294967295.
- **Route Map**—Checks for filtering the importing of routes from the source routing protocol to the current routing protocol. If this parameter is not specified, all routes are redistributed. If this parameter is specified, but no route map tags are listed, no routes are imported. Or you can add a new route map by clicking **Add** (+). See Route Map, on page 1426 for the procedure to add a new route map.
- **Process ID**—The OSPF process ID, either 1 or 2.

  Note
  The Process ID is enabled the OSPFv3 process is redistributing a route learned by another OSPFv3 process.

- **Match**—Enables OSPF routes to be redistributed into other routing domains:
  - **Internal** for routes that are internal to a specific autonomous system.
  - **External 1** for routes that are external to the autonomous system, but are imported into OSPFv3 as Type 1 external routes.
  - **External 2** for routes that are external to the autonomous system, but are imported into OSPFv3 as Type 2 external routes.
  - **NSSA External 1** for routes that are external to the autonomous system, but are imported into OSPFv3 in an NSSA for IPv6 as Type 1 external routes.
  - **NSSA External 2** for routes that are external to the autonomous system, but are imported into OSPFv3 in an NSSA for IPv6 as Type 2 external routes.

Step 5 Click **OK** to save the redistribution configuration.

Step 6 Click **Save** on the Routing page to save your changes.

What to do next

Continue with Configure OSPFv3 Summary Prefixes, on page 1260.

Configure OSPFv3 Summary Prefixes

You can configure the Firewall Threat Defense device to advertise routes that match a specified IPv6 prefix and mask pair.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select **Routing** > **OSPFv3**.

Step 3 Select **Summary Prefix** > **Add**.

You can click **Edit** (), or use the right-click menu to cut, copy, past, insert, and delete summary prefixes.

Step 4 Configure the following summary prefix options for each OSPF process:

- **IPv6 Prefix/Length**—The IPv6 prefix and prefix length label. Select one from the list or click **Add** () to add a new network object. See Network, on page 1401 for the procedure for adding networks.
- **Advertise**— Advertises routes that match the specified prefix and mask pair. Uncheck this check box to suppress routes that match the specified prefix and mask pair.
- (Optional) **Tag**—A value that you can use as a match value for controlling redistribution through route maps.

Step 5 Click **OK** to save the summary prefix configuration.

Step 6 Click **Save** on the Routing page to save your changes.

What to do next

Continue with Configure OSPFv3 Interfaces, Authentication, and Neighbors, on page 1261.

Configure OSPFv3 Interfaces, Authentication, and Neighbors

You can change certain interface-specific OSPFv3 parameters, if necessary. You are not required to change any of these parameters, but the following interface parameters must be consistent across all routers in an attached network: the hello interval and the dead interval. If you configure any of these parameters, be sure that the configurations for all routers on your network have compatible values.

For a successful implementation of OSPFv3 authentication with Nexus switches, ensure that you have the compatible version of the switch, for example, Nexus 3000, 7000, and 9000 series switches.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select **Routing** > **OSPFv3**.

Step 3 Select **Interface** > **Add**.

You can click **Edit** to edit, or use the right-click menu to cut, copy, past, insert, and delete areas.

Step 4 Configure the following interface options for each OSPFv3 process:

- **Interface**—The interface you are configuring.

- **Enable OSPFv3**—Enables OSPFv3.
- **OSPF Process**—Choose 1 or 2.
- **Area**—The area ID for this process.
- **Instance** —Specifies the area instance ID to be assigned to the interface. An interface can have only one OSPFv3 area. You can use the same area on multiple interfaces, and each interface can use a different area instance ID.

Step 5 Select **Properties**, and configuring the following options for each OSPFv3 process:

- **Filter Outgoing Link Status Advertisements**—Filters outgoing LSAs to an OSPFv3 interface. All outgoing LSAs are flooded to the interface by default.
- **Disable MTU mismatch detection**—Disables the OSPF MTU mismatch detection when DBD packets are received. OSPF MTU mismatch detection is enabled by default.
- **Flood Reduction**—Changes normal LSAs into Do Not Age LSAs, so that they don't get flooded every 3600 seconds across areas.

  OSPF LSAs are refreshed every 3600 seconds. In large OSPF networks, this can lead to large amounts of unnecessary LSA flooding from area to area.
- **Point-to-Point Network**—Lets you transmit OSPF routes over VPN tunnels. When an interface is configured as point-to-point, non-broadcast, the following restrictions apply:
  - You can define only one neighbor for the interface.
  - You need to manually configure the neighbor.
  - You need to define a static route pointing to the crypto endpoint.
  - If OSPF over a tunnel is running on the interface, regular OSPF with an upstream router cannot be run on the same interface.
  - You should bind the crypto map to the interface before specifying the OSPF neighbor to ensure that the OSPF updates are passed through the VPN tunnel. If you bind the crypto map to the interface after specifying the OSPF neighbor, use the **clear local-host all** command to clear OSPF connections so that the OSPF adjacencies can be established over the VPN tunnel.
- **Broadcast**— Specifies that the interface is a broadcast interface. By default, this check box is checked for Ethernet interfaces. Uncheck this check box to designate the interface as a point-to-point, nonbroadcast interface. Specifying an interface as point-to-point, nonbroadcast lets you transmit OSPF routes over VPN tunnels.
- **Cost**—Specifies the cost of sending a packet on the interface. Valid values for this setting range from 0 to 255. The default value is 1. Entering 0 for this setting makes the router ineligible to become the designated router or backup designated router. This setting does not apply to interfaces that are configured as point-to-point, nonbroadcast interfaces.

  When two routers connect to a network, both attempt to become the designated router. The device with the higher router priority becomes the designated router. If there is a tie, the router with the higher router ID becomes the designated router.
- **Priority**—Determines the designated router for a network. Valid values range from 0 to 255.

- **Dead Interval**—Time period in seconds for which hello packets must not be seen before neighbors indicate that the router is down. The value must be the same for all nodes on the network and can range from 1 to 65535.
- **Hello Interval**— Time period in seconds between OSPF packets that the router will send before adjacency is established with a neighbor. Once the routing device detects an active neighbor, the hello packet interval changes from the time specified in the poll interval to the time specified in the hello interval. Valid values range from 1 to 65535 seconds.
- **Retransmit Interval**—Time in seconds between LSA retransmissions for adjacencies that belong to the interface. The time must be greater than the expected round-trip delay between any two routers on the attached network. Valid values range from 1 to 65535 seconds. The default is 5 seconds.
- **Transmit Delay**—Estimated time in seconds to send a link-state update packet on the interface. Valid values range from 1 to 65535 seconds. The default is 1 second.
- **Enable BFD**—Allows you to enable BFD on this interface.

Step 6 Click **OK** to save the properties configuration.

Step 7 Select **Authentication**, and configure the following options for each OSPFv3 process:

- **Type**—Type of authentication. The available options are Area, Interface, and None. The None option indicates that no authentication is used.
- **Security Parameters Index**— A number from 256 to 4294967295. Configure this if you chose Interface as the type.
- **Authentication**—Type of authentication algorithm. Supported values are SHA-1 and MD5. Configure this if you chose Interface as the type.
- **Authentication Key**— When MD5 authentication is used, the key must be 32 hexadecimal digits (16 bytes) long. When SHA-1 authentication is used, the key must be 40 hexadecimal digits (20 bytes) long.
- **Encrypt Authentication Key**—Enables encryption of the authentication key.
- **Include Encryption**— Enables encryption.
- **Encryption Algorithm**—Type of encryption algorithm. Supported value is DES. The NULL entry indicates no encryption. Configure this if you chose **Include Encryption**.
- **Encryption Key**—Enter the encryption key. Configure this if you chose **Include Encryption**.
- **Encrypt Key**—Enables the key to be encrypted.

Step 8 Click **OK** to save the authentication configuration.

Step 9 Select **Neighbor**, click **Add**, and configure the following options for each OSPFv3 process:

- **Link Local Address**—The IPv6 address of the static neighbor.
- **Cost**—Enables cost. Enter the cost in the **Cost** field, and check the **Filter Outgoing Link State Advertisements** if you want to advertise.
- (Optional) **Poll Interval**—Enables the poll interval. Enter the **Priority** level and the **Poll Interval** in seconds.

Step 10 Click **Add** to add the neighbor.

Step 11 Click **OK** to save the Interface configuration.

Configure OSPFv3 Advanced Properties

The Advanced Properties allows you to configure options, such as syslog message generation, administrative route distances, passive OSPFv3 routing, LSA timers, and graceful restarts.

Graceful Restarts

The Firewall Threat Defense device may experience some known failure situations that should not affect packet forwarding across the switching platform. The Non-Stop Forwarding (NSF) capability allows data forwarding to continue along known routes, while the routing protocol information is being restored. This capability is useful when there is a scheduled hitless software upgrade. You can configure graceful restart on OSPFv3 using graceful-restart (RFC 5187).

Note NSF capability is also useful in HA mode and clustering.

Configuring the NSF graceful-restart feature involves two steps; configuring capabilities and configuring a device as NSF-capable or NSF-aware. A NSF-capable device can indicate its own restart activities to neighbors and a NSF-aware device can help a restarting neighbor.

A device can be configured as NSF-capable or NSF-aware, depending on some conditions:

- A device can be configured as NSF-aware irrespective of the mode in which it is.
- A device has to be in either Failover or Spanned Etherchannel (L2) cluster mode to be configured as NSF-capable.
- For a device to be either NSF-aware or NSF-capable, it should be configured with the capability of handling opaque Link State Advertisements (LSAs)/ Link Local Signaling (LLS) block as required.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **OSPFv3** > **Advanced**.

Step 3 For **Router ID**, choose Automatic or IP Address(appears for non-cluster and a cluster in spanned etherchannel mode) or Cluster Pool (appears for a cluster in individual interface mode). If you choose IP Address, enter the IPv6 address in the **IP Address** field. If you choose Cluster Pool, choose the IPv6 cluster pool value from the **Cluster Pool** down-down field. For information on creating the cluster pool address, see Address Pools, on page 1374.

Step 4 Check the **Ignore LSA MOSPF** check box if you want to suppress syslog messages when the route receives unsupported LSA Type 6 multicast OSPF (MOSPF) packets.

Step 5 Select **General**, and configure the following:

- **Adjacency Changes**—Defines the adjacency changes that cause syslog messages to be sent.

  By default, a syslog message is generated when an OSPF neighbor goes up or down. You can configure the router to send a syslog message when an OSPF neighbor goes down and also a syslog for each state.

- **Adjacency Changes**—Causes the Firewall Threat Defense device to send a syslog message whenever an OSPF neighbor goes up or down. This setting is checked by default.
- **Include Details**—Causes the Firewall Threat Defense device to send a syslog message whenever any state change occurs, not just when a neighbor goes up or down. This setting is unchecked by default.

- **Administrative Route Distances**—Allows you to modify the settings that were used to configure administrative route distances for inter-area, intra-area, and external IPv6 routes. The administrative route distance is an integer from 1 to 254. The default is 110.
- **Default Information Originate**—Check the **Enable** check box to generate a default external route into an OSPFv3 routing domain and configure the following options:
  - **Always Advertise**—Will always advertise the default route whether or not one exists.
  - **Metric**—Metric used for generating the default route. Valid metric values range from 0 to 16777214. The default value is 10.
  - **Metric Type**—The external link type that is associated with the default route that is advertised into the OSPFv3 routing domain. Valid values are 1 (Type 1 external route) and 2 (Type 2 external route). The default is Type 2 external route.
  - **Route Map**—Choose the routing process that generates the default route if the route map is satisfied or click **Add** (+) to add a new one. See Route Map, on page 1426 to add a new route map.

Step 6 Click **OK** to save the general configuration.

Step 7 Select **Passive Interface**, select the interfaces on which you want to enable passive OSPFv3 routing from the Available Interfaces list, and click **Add** to move them to the Selected Interfaces list.

Passive routing assists in controlling the advertisement of OSPFv3 routing information and disables the sending and receiving of OSPFv3 routing updates on an interface.

Step 8 Click **OK** to save the passive interface configuration.

Step 9 Select **Timer**, and configure the following LSA pacing and SPF calculation timers:

- **Arrival**—Specifies the minimum delay in milliseconds that must pass between acceptance of the same LSA arriving from neighbors. The range is from 0 to 6000,000 milliseconds. The default is 1000 milliseconds.
- **Flood Pacing**—Specifies the time in milliseconds at which LSAs in the flooding queue are paced in between updates. The configurable range is from 5 to 100 milliseconds. The default value is 33 milliseconds.
- **Group Pacing**—Specifies the interval in seconds at which LSAs are collected into a group and refreshed, check summed, or aged. Valid values range from 10 to 1800. The default value is 240.
- **Retransmission Pacing**—Specifies the time in milliseconds at which LSAs in the retransmission queue are paced. The configurable range is from 5 to 200 milliseconds. The default value is 66 milliseconds.
- **LSA Throttle**—Specifics the delay in milliseconds to generate the first occurrence of the LSA. The default value is 0 millisecond. The minimum specifies the minimum delay in milliseconds to originate the same LSA. The default value is 5000 milliseconds. The maximum specifies the maximum delay in milliseconds to originate the same LSA. The default value is 5000 milliseconds.

Note
For LSA throttling, if the minimum or maximum time is less than the first occurrence value, then OSPFv3 automatically corrects to the first occurrence value. Similarly, if the maximum delay specified is less than the minimum delay, then OSPFv3 automatically corrects to the minimum delay value.

- **SPF Throttle**—Specifies the delay in milliseconds to receive a change to the SPF calculation. The default value is 5000 milliseconds. The minimum specifies the delay in milliseconds between the first and second SPF calculations. The default value is 10000 milliseconds. The maximum specifies the maximum wait time in milliseconds for SPF calculations. The default value is 10000 milliseconds.

  Note
  For SPF throttling, if the minimum or maximum time is less than the first occurrence value, then OSPFv3 automatically corrects to the first occurrence value. Similarly, if the maximum delay specified is less than the minimum delay, then OSPFv3 automatically corrects to the minimum delay value.

Step 10 Click **OK** to save the LSA timer configuration.

Step 11 Select **Non Stop Forwarding**, and check the **Enable graceful-restart helper** check box. This is checked by default. Uncheck this to disable the graceful-restart helper mode on an NSF-aware device.

Step 12 Check the **Enable link state advertisement** check box to enable strict link state advertisement checking.

When enabled, it indicates that the helper router will terminate the process of restarting the router if it detects that there is a change to a LSA that would be flooded to the restarting router, or if there is a changed LSA on the retransmission list of the restarting router when the graceful restart process is initiated.

Step 13 Check the **Enable graceful-restart (Use when Spanned Cluster or Failover Configured)** and enter the graceful-restart interval in seconds. The range is 1-1800. The default value is 120 seconds. For a restart interval below 30 seconds, graceful restart will be terminated.

Step 14 Click **OK** to save the graceful restart configuration.

Step 15 Click **Save** on the Routing page to save your changes.

History for OSPF

Table 71: Feature History for OSPF

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
BFD Support for OSPF v2 and v3	7.4	7.4	You can enable BFD on OSPFv2 and OSPFv3 interfaces. New/modified screens: • **Configuration** > **Device Setup** > **Routing** > **OSPFv2** • **Configuration** > **Device Setup** > **Routing** > **OSPFv3**

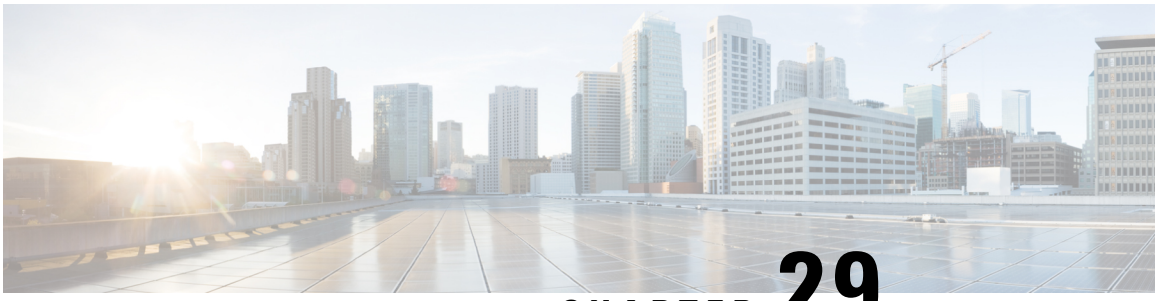

CHAPTER 29

EIGRP

This section describes how to configure the Firewall Threat Defense to route data, perform authentication, and redistribute routing information using the Enhanced Interior Gateway Routing Protocol (EIGRP).

- About EIGRP Routing, on page 1267
- Requirements and Prerequisites for EIGRP, on page 1268
- Guidelines and Limitations of EIGRP Routing, on page 1268
- Configure EIGRP, on page 1270
- History for EIGRP, on page 1276

About EIGRP Routing

Enhanced Interior Gateway Routing Protocol (EIGRP), developed by Cisco, is an enhanced version of IGRP. Unlike IGRP and RIP, EIGRP does not send out periodic route updates. EIGRP updates are sent out only when the network topology changes. Key capabilities that distinguish EIGRP from other routing protocols include fast convergence, support for variable-length subnet mask, support for partial updates, and support for multiple network layer protocols.

A router running EIGRP stores all the neighbor routing tables so that it can quickly adapt to alternate routes. If no appropriate route exists, EIGRP queries its neighbors to discover an alternate route. These queries are propagated until an alternate route is found. EIGRP support for the variable-length subnet masks allows routes to be automatically summarized on a network boundary. Additionally, EIGRP can be configured to summarize any bit boundary at any interface.

EIGRP does not make periodic updates. Instead, it sends partial updates when the metric for a route changes. Propagation of partial updates is automatically bounded such that only those routers that need the information are updated. As a result of these two capabilities, EIGRP consumes significantly less bandwidth than IGRP.

To dynamically learn of other routers on directly attached networks, threat defense uses neighbor discovery. EIGRP routers send out multicast hello packets to announce their presence on the network. When the EIGRP device receives a hello packet from a new neighbor, it sends its topology table to the neighbor with an initialization bit set. When the neighbor receives the topology update with the initialization bit set, the neighbor sends its topology table back to the device.

The hello packets are sent out as multicast messages. No response is expected for a hello message. Statically defined neighbors is an exception to this rule. If you manually configure a neighbor, hello messages, routing updates, and acknowledgments are sent as unicast messages.

Once this neighbor relationship is established, routing updates are not exchanged unless there is a change in the network topology. The neighbor relationship is maintained through the hello packets. Each hello packet

received from a neighbor includes a hold time. Hold time is the time within which threat defense can expect to receive a hello packet from that neighbor. If the device does not receive a hello packet from that neighbor within the hold time advertised by that neighbor, the device considers that neighbor to be unavailable.

EIGRP uses neighbor discovery/recovery, Reliable Transport Protocol (RTP), and Diffusing Update Algorithm (DUAL) for route computations. DUAL saves all routes to a destination in the topology table, and not just the least-cost route. The least-cost route is inserted into the routing table. The other routes remain in the topology table. If the main route fails, another route is chosen from the feasible successors. A successor is a neighboring router that is used for packet forwarding that has a least-cost path to a destination. A feasibility calculation ensures that the path is not part of a routing loop.

If a feasible successor is not found in the topology table, a route recomputation takes place. During route recomputation, DUAL queries the EIGRP neighbors for a route. The query is propagated to successive neighbors. If a feasible successor is not found, an unreachable message is returned.

During route recomputation, DUAL marks the route as active. By default, threat defense waits for three minutes to receive a response from its neighbors. If the device does not receive a response from a neighbor, the route is marked as stuck-in-active. All routes in the topology table that point to the unresponsive neighbor as a feasibility successor are removed.

Requirements and Prerequisites for EIGRP

Model Support

Threat Defense

Firewall Threat Defense Virtual

Supported Domains

Any

User Roles

Admin

Network Admin

Guidelines and Limitations of EIGRP Routing

Firewall Mode Guidelines

Supported on routed firewall mode only.

Device Guidelines

- Only one EIGRP process is allowed per device.
- EIGRP can be configured through management center UI on Firewall Threat Defense 6.6 and higher versions.

Interface Guidelines

- Only routed interfaces with logical names and with an IP address can be associated with an EIGRP routing process.
- Only interfaces belonging to the global virtual router can be part of EIGRP. EIGRP can learn, filter, and redistribute routes across routing protocols in global virtual router.
- Supports physical, EtherChannel, redundant, subinterfaces only. However, the members of EtherChannel interfaces are not supported.
- BVI and VNI cannot be part of EIGRP.
- A passive interface cannot be configured as a neighbor interface.

IP Address and Network Objects Support

- Only IPv4 address is supported.
- Range, FQDN, and wildcard mask are not supported.
- Only Standard access list objects are supported.

Redistribution Guidelines

- BGP, OSPF, and RIP in the global virtual router can redistribute to EIGRP.
- EIGRP can redistribute to BGP, OSPF, RIP, Static, and Connected in the global virtual router.
- When EIGRP is configured on a device that is a part of OSPF network or vice versa, ensure that OSPF-router is configured to tag the route (EIGRP does not support route tag).

  When redistributing EIGRP into OSPF and OSPF into EIGRP, a routing loop occurs when there is an outage on one of the links, interfaces, or even when the route originator is down. To prevent the redistribution of routes from one domain back into the same domain, a router can tag a route that belongs to a domain while it is redistributing, and those routes can be filtered on the remote router based on the same tag. Because the routes will not be installed into the routing table, they will not be redistributed back into the same domain.

Deployment Process Guidelines

When you want to change the existing AS number of a deployed EIGRP configuration, you must disable the EIGRP and deploy it. This step will clear the deployed EIGRP configuration on the threat defense. Next, recreate the EIGRP configurations with a new AS number and then deploy it. Thus, this process prevents any deployment failures owing to the same EIGRP configuration being deployed on the threat defense.

Upgrade Guidelines

When you upgrade to version 7.2 and later when the previous version has any FlexConfig EIGRP policies, the management center displays a warning message during deployment. However, it does not stop the deployment process. However, after deployment, to manage the EIGRP policies from the UI (**Device (Edit)** > **Routing** > **EIGRP**), you must redo the configuration in the **Device (Edit)** > **Routing** > **EIGRP** page and remove the configuration from FlexConfig. To automate creation of the policies in the UI, Firewall Management Center provides an option to migrate the policies from FlexConfig to the UI. For more details, see Migrating FlexConfig Policies, on page 2708.

Configure EIGRP

You can enable and configure EIGRP on the firewall device in the **Routing** tab.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click the **Routing** tab.

Step 3 Under Global, click **EIGRP**.

Step 4 Check the **Enable EIGRP** check box to enable the EIGRP routing process.

Step 5 In the **AS Number** field, enter the autonomous system (AS) number for the EIGRP process. The AS number includes multiple autonomous numbers. The AS number can be from 1 to 65535 and is a uniquely assigned value that identifies each network on the Internet.

Step 6 To configure other EIGRP properties, see the following topics:

a. Configure EIGRP Settings, on page 1270.

b. Configure EIGRP Neighbors Settings, on page 1271.

c. Configure EIGRP Filter Rules Settings, on page 1271.

d. Configure EIGRP Redistribution Settings, on page 1272.

e. Configure EIGRP Summary Address Settings, on page 1273.

f. Configure EIGRP Interfaces Settings, on page 1274.

g. Configure EIGRP Advanced Settings, on page 1274.

Configure EIGRP Settings

Procedure

Step 1 On the **EIGRP** page, click the **Setup** tab.

Step 2 Check the **Auto Summary** check box to enable EIGRP to summarize network number boundaries.

Note
Enabling Auto Summary can cause routing problems if you have noncontiguous networks.

Step 3 In the **Available Networks/Hosts** box, click the networks or hosts that should participate in the EIGRP routing process, and then click **Add**. To add a new network object, click **Add** (+). See Network, on page 1401 for the procedure for adding networks.

Step 4 To configure passive interfaces, check the **Passive Interface** check box. In EIGRP, a passive interface does not send or receive routing updates.

a) To specify selective interfaces as passive, click the **Selected Interface** radio button. In the **Available Interfaces** box, select the interfaces, and click **Add**.
b) To specify all interfaces as passive, click the **All Interfaces** radio button.

Step 5 Click **Ok** and **Save** the settings.

Configure EIGRP Neighbors Settings

You can define static neighbors for the EIGRP process. When you define an EIGRP neighbor, hello packets are unicast to that neighbor.

Procedure

Step 1 On the **EIGRP** page, click the **Neighbors** tab.
Step 2 Click **Add**.
Step 3 From the **Interface** drop-down, choose the interface through which the neighbor is available.
Step 4 From the **Neighbor** drop-down, choose the IP address of the static neighbor. To add the network object, click **Add** (+). See Network, on page 1401 for the procedure for adding network objects.
Step 5 Click **Ok** and **Save** the settings.

Configure EIGRP Filter Rules Settings

You can configure route filtering rules for the EIGRP routing process. Filter rules allow you to control the routes that are accepted or advertised by the EIGRP routing process.

Procedure

Step 1 On the **EIGRP** page, click the **Filter Rules** tab.
Step 2 Click **Add** (+).
Step 3 In the **Add Filter Rules** dialog box, from the **Filter Direction** drop-down, choose the direction for the rule:

- Inbound—The rule filters default route information from incoming EIGRP routing updates.
- Outbound—The rule filters default route information from outgoing EIGRP routing updates.

Step 4 To select the interface to which the filtering rule applies, click the **Interface** radio button, and from the drop-down, select the interface.

Note
You cannot apply EIGRP filtering rules on VTI interfaces.

Step 5 To select the protocol to which the filtering rule applies, click the **Protocol** radio button and from the drop-down, select the protocol—BGP, RIP, Static, Connected, or OSPF. For BGP and OSPF protocols, you can specify the relevant Process ID.

Step 6 From the **Access List** drop-down, choose the access list. The list defines the networks that are to be received and that are to be suppressed in routing updates. To add a new standard access list object, click **Add** (+) and see Configure Standard ACL Objects, on page 1374 for the detailed procedure.

Step 7 Click **Ok** and **Save** the settings.

Configure EIGRP Redistribution Settings

You can define the rules for redistributing routes from other routing protocols to the EIGRP routing process.

Procedure

Step 1 On the **EIGRP** page, click the **Redistribution** tab.

Step 2 Click **Add** (+).

Step 3 In the **Add Redistribution** dialog box, from the **Protocol** drop-down, choose the source protocol from which the routes are being redistributed:

- BGP—Redistributes routes discovered by the BGP routing process to EIGRP.
- RIP—Redistributes routes discovered by the RIP routing process to EIGRP.
- Static—Redistributes static routes to the EIGRP routing process. Static routes that fall within the scope of a network statement are automatically redistributed to EIGRP; you do not need to define a redistribution rule for them. However, when redistributing static routes that point to VTI interfaces in EIGRP, you must specify the metric. For static routes pointing to other types of interfaces, specifying the metric is not required.
- Connected—Redistributes connected routes (routes established automatically by virtue of having IP address enabled on the interface) to the EIGRP routing process. Connected routes that fall within the scope of a network statement are automatically redistributed to EIGRP; you do not need to define a redistribution rule for them.
- OSPF—Redistributes routes discovered by the OSPF routing process to EIGRP. If you choose this protocol, the Match options on this dialog box become available under **Optional OSPF Redistribution**:
  - Internal—Routes that are internal to a specific AS.
  - External1—Routes that are external to the AS and imported into OSPF as a Type 1 external route.
  - External2—Routes that are external to the AS and imported into the selected process as a Type 2 external route.
  - Nsaa-External1—Not-So-Stubby Area (NSSA) routes that are external to the AS and imported into the selected process as Type 1 external routes.
  - Nsaa-External2—(NSSA) routes that are external to the AS and imported into the selected process as Type 2 external routes.

Note
These options are not available when redistributing static, connected, RIP, or BGP routes.

Step 4 Under **Optional Metrics** enter the relevant values:

- **Bandwidth**—The minimum bandwidth of the route in kilobits per second. Valid values range from 1 to 4294967295.
- **Delay Time**—The routing delay in tens of microseconds. Valid values range from 0 to 4294967295.
- **Reliability** —The likelihood of successful packet transmission is expressed as a number 0 through 255. The value 255 indicates 100 percent reliability; 0 means no reliability.
- **Loading**— The effective bandwidth of the route. Valid values range from 1 to 255. 255 indicates 100 percent loading.
- **MTU**—The smallest permissible value for the maximum transmission unit of the path. Valid values range from 1 to 65535.

Step 5 From the **Route Map** drop-down, choose the route map object to apply to the redistribution entry. To create a new route map object, click **Add** (+). See Configure Route Map Entry for the procedure to add a new route map.

Step 6 Click **Ok** and **Save** the settings.

Configure EIGRP Summary Address Settings

You can configure summary addresses for each interface. You need to manually define summary addresses if you want to create summary addresses that do not occur at a network boundary, or if you want to use summary addresses on threat defense with automatic route summarization disabled. If more specific routes are available in the routing table, EIGRP advertises the summary address with a metric equal to the minimum of all the more specific routes.

Procedure

Step 1 On the **EIGRP** page, click the **Summary Address** tab.

Step 2 Click **Add**.

Step 3 From the **Interface** drop-down, choose the interface from which the summary address is advertised.

Step 4 From the **Network** drop-down, choose the network object with specific IP address and network mask to be summarized. To add a new network, click **Add** (+). See Network, on page 1401 for the procedure for adding networks.

Step 5 In the **Administrative Distance** field, enter the administrative distance of the summary route. Valid values range from 1 to 255.

Step 6 Click **Ok** and **Save** the settings.

Configure EIGRP Interfaces Settings

You can configure interface-specific EIGRP routing properties in the Interfaces tab.

Procedure

Step 1 On the **EIGRP** page, click the **Interfaces** tab.

Step 2 Click **Add** (+).

Step 3 From the **Interface** drop-down, choose the name of the interface to which the configuration applies.

Step 4 In the **Hello Interval** field, enter the interval, in seconds, between EIGRP hello packets that are sent on an interface. Valid values range from 1 to 65535. The default value is 5 seconds.

Step 5 In the **Hold Time** field, enter the hold time that is advertised by the device in EIGRP hello packets. Valid values range from 3 to 65535. The default value is 15 seconds.

Step 6 To enable EIGRP split-horizon on the interface, click the **Split Horizon** check box.

Step 7 In the **Delay Time** field, enter the delay time in tens of microseconds. Valid values are from 1 to 16777215. This option is not supported for devices in multi-context mode.

Step 8 Specify values for the Authentication properties:

- **Enable MD5 Authentication**—Check the check box to use MD5 hash algorithm for authentication of EIGRP packets.
- **Key Type**—From the drop-down, select any one of the following key type:
  - None—To indicate that no authentication is required.
  - Unencrypted—To indicate that the key string to be used is a clear text password for authentication.
  - Encrypted—To indicate that the key string to be used is an encrypted password for authentication.
  - Auth Key—To indicate that the key string to be used is an EIGRP authentication key.
- **Key ID**—The ID of the key that is used to authenticate EIGRP updates. Enter a numerical key identifier. Valid values range from 0 to 255.
- **Key**—An alphanumeric character string of up to 17 characters. For an encrypted authentication type, this field should have a minimum of 17 characters.
- **Confirm Key**—Re-enter the key.

Step 9 Click **Ok** and **Save** the settings.

Configure EIGRP Advanced Settings

You can configure EIGRP advanced settings such as the router ID, stub routing, and adjacency changes.

Procedure

Step 1 On the **EIGRP** page, click the **Advanced** tab.

Step 2 Under **Default Route Information**, you can specify the sending and receiving of default route information in EIGRP updates.

- (Appears for non-cluster and cluster in spanned etherchannel mode)**Router ID (IP Address)**—Enter the ID used to identify the originating router for external routes. If an external route is received with the local router ID, the route is discarded. To prevent this issue, specify a global address for the router ID. An unique value should be configured for each EIGRP router.
- (Appears only for a cluster in individual interface mode)**IPv4 Address Pool**—Select the relevant cluster pool value (IPv4 address pool object). To create the address pool, see Address Pools, on page 1374.
- **Accept Default Route Info**—Check the check box to configure EIGRP to accept exterior default routing information.
  - **Access List**—From the **Access List** drop-down, specify a standard access list that defines the networks that are allowed and the networks that are not when receiving default route information.

    To add a new standard access list object, click **Add** (+) and see Configure Standard ACL Objects, on page 1374 for the detailed procedure.
- **Send Default Route Info**—Check the check box to configure EIGRP to advertise exterior default routing information.
  - **Access List**—From the **Access List** drop-down, specify a standard access list that defines the networks that are allowed and the networks that are not when sending default route information.

    To add a new standard access list object, click **Add** (+) and see Configure Standard ACL Objects, on page 1374 for the detailed procedure.

Step 3 Under **Administrative Distance**, specify:

- **Internal Distance**—Administrative distance for EIGRP internal routes. Internal routes are those that are learned from another entity within the same autonomous system. Valid values range from 1 to 255. The default value is 90.
- **External Distance**—Administrative distance for EIGRP external routes. External routes are those for which the best path is learned from a neighbor external to the autonomous system. Valid values range from 1 to 255. The default value is 170.

Step 4 Under **Adjacency Changes**, specify:

- **Log Neighbor Changes**—Click the check box to enable the logging of EIGRP neighbor adjacency changes.
- **Log Neighbor Warnings**—Click the check box to enable the logging of EIGRP neighbor warning messages.
- (Optional) Enter the time interval (in seconds) between repeated neighbor warning messages. Valid values range from 1 to 65535. Repeated warnings are not logged if they occur during this interval.

Step 5 Under **Stub**, to enable the device as an EIGRP stub routing process, click one or more of the following EIGRP stub routing processes check boxes:

- **Receive only**—Configures the EIGRP stub routing process to receive route information from the neighbor routers but not send route information to the neighbors. If this option is selected, you cannot select any of the other stub routing options.
- **Connected**—Advertises connected routes.
- **Redistributed**—Advertises redistributed routes.
- **Static**—Advertises static routes.
- **Summary**—Advertises summary routes.

Step 6 Under **Default Metrics**, define the default metrics for routes redistributed to the EIGRP routing process:

- **Bandwidth**—the minimum bandwidth of the route in kilobits per second. Valid values range from 1 to 4294967295.
- **Delay Time**—the route delay in tens of microseconds. Valid values range from 0 to 4294967295.
- **Reliability**—the likelihood of successful packet transmission expressed as a number 0 through 255. The value 255 indicates 100 percent reliability; 0 means no reliability.
- **Loading**—the effective bandwidth of the route. Valid values range from 1 to 255; 255 indicates 100 percent loading.
- **MTU**—the smallest allowed value for the maximum transmission unit of the path. Valid values range from 1 to 65535.

History for EIGRP

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
EIGRP configuration	7.2	Any	In the previous releases, EIGRP was configurable on threat defense only through FlexConfig. FlexConfig no longer supports EIGRP configuration. You can now configure EIGRP settings for threat defense in the management center UI. New/modified screens: **Devices** > **Device Management** > **Routing** > **EIGRP**.

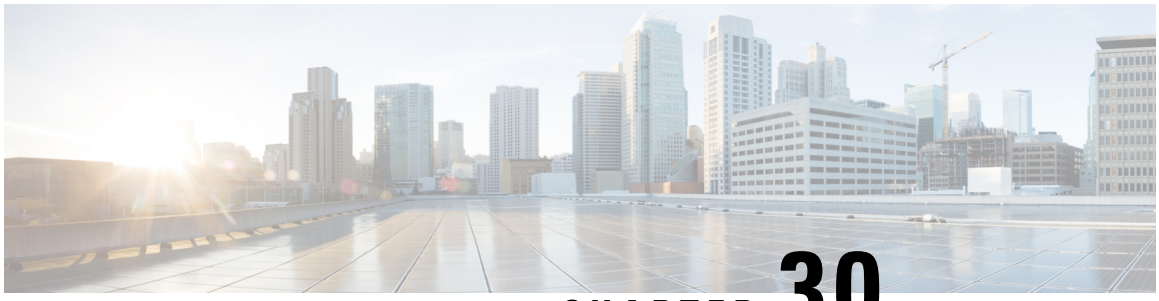

CHAPTER 30

BGP

This section describes how to configure the Firewall Threat Defense to route data, perform authentication, and redistribute routing information using the Border Gateway Protocol (BGP).

- About BGP, on page 1277
- Requirements and Prerequisites for BGP, on page 1280
- Guidelines for BGP, on page 1280
- Configure BGP, on page 1281
- History for BGP, on page 1295

About BGP

BGP is an inter and intra autonomous system routing protocol. An autonomous system is a network or group of networks under a common administration and with common routing policies. BGP is used to exchange routing information for the Internet and is the protocol used between Internet service providers (ISP).

Routing Table Changes

BGP neighbors exchange full routing information when the TCP connection between neighbors is first established. When changes to the routing table are detected, the BGP routers send to their neighbors only those routes that have changed. BGP routers do not send periodic routing updates, and BGP routing updates advertise only the optimal path to a destination network.

Note AS loop detection is done by scanning the full AS path (as specified in the AS_PATH attribute), and checking that the AS number of the local system does not appear in the AS path. By default, EBGP advertises the learned routes to the same peer to prevent additional CPU cycles on the device in performing loop checks and to avoid delays in the existing outgoing update tasks.

Routes learned via BGP have properties that are used to determine the best route to a destination, when multiple paths exist to a particular destination. These properties are referred to as BGP attributes and are used in the route selection process:

- Weight—This is a Cisco-defined attribute that is local to a router. The weight attribute is not advertised to neighboring routers. If the router learns about more than one route to the same destination, the route with the highest weight is preferred.

- Local preference—The local preference attribute is used to select an exit point from the local AS. Unlike the weight attribute, the local preference attribute is propagated throughout the local AS. If there are multiple exit points from the AS, the exit point with the highest local preference attribute is used as an exit point for a specific route.
- Multi-exit discriminator—The multi-exit discriminator (MED) or metric attribute is used as a suggestion to an external AS regarding the preferred route into the AS that is advertising the metric. It is referred to as a suggestion because the external AS that is receiving the MEDs may also be using other BGP attributes for route selection. The route with the lower MED metric is preferred.
- Origin—The origin attribute indicates how BGP learned about a particular route. The origin attribute can have one of three possible values and is used in route selection.
  - IGP—The route is interior to the originating AS. This value is set when the network router configuration command is used to inject the route into BGP.
  - EGP—The route is learned via the Exterior Border Gateway Protocol (EBGP).
  - Incomplete—The origin of the route is unknown or learned in some other way. An origin of incomplete occurs when a route is redistributed into BGP.
- AS_path—When a route advertisement passes through an autonomous system, the AS number is added to an ordered list of AS numbers that the route advertisement has traversed. Only the route with the shortest AS_path list is installed in the IP routing table.
- Next hop—The EBGP next-hop attribute is the IP address that is used to reach the advertising router. For EBGP peers, the next-hop address is the IP address of the connection between the peers. For IBGP, the EBGP next-hop address is carried into the local AS. However, when the next hop is in the same subnet as the peering address of the eBGP peer, the next hop is not modified. This behavior is referred to as the third party next hop.

  Use the **next-hop-self** command when redistributing VPN-advertised routes to iBGP peers to ensure that the routes are redistributed with the correct next hop IP.
- Community—The community attribute provides a way of grouping destinations, called communities, to which routing decisions (such as acceptance, preference, and redistribution) can be applied. Route maps are used to set the community attribute. The predefined community attributes are as follows:
  - no-export—Do not advertise this route to EBGP peers.
  - no-advertise—Do not advertise this route to any peer.
  - internet—Advertise this route to the Internet community; all routers in the network belong to it.

When to Use BGP

Customer networks, such as universities and corporations, usually employ an Interior Gateway Protocol (IGP) such as OSPF for the exchange of routing information within their networks. Customers connect to ISPs, and ISPs use BGP to exchange customer and ISP routes. When BGP is used between autonomous systems (AS), the protocol is referred to as External BGP (EBGP). If a service provider is using BGP to exchange routes within an AS, then the protocol is referred to as Interior BGP (IBGP).

BGP can also be used for carrying routing information for IPv6 prefix over IPv6 networks.

BGP Path Selection

BGP may receive multiple advertisements for the same route from different sources. BGP selects only one path as the best path. When this path is selected, BGP puts the selected path in the IP routing table and propagates the path to its neighbors. BGP uses the following criteria, in the order presented, to select a path for a destination:

- If the path specifies a next hop that is inaccessible, drop the update.
- Prefer the path with the largest weight.
- If the weights are the same, prefer the path with the largest local preference.
- If the local preferences are the same, prefer the path that was originated by BGP running on this router.
- If no route was originated, prefer the route that has the shortest AS_path.
- If all paths have the same AS_path length, prefer the path with the lowest origin type (where IGP is lower than EGP, and EGP is lower than incomplete).
- If the origin codes are the same, prefer the path with the lowest MED attribute.
- If the paths have the same MED, prefer the external path over the internal path.
- If the paths are still the same, prefer the path through the closest IGP neighbor.
- Determine if multiple paths require installation in the routing table for BGP Multipath, on page 1279.
- If both paths are external, prefer the path that was received first (the oldest one).
- Prefer the path with the lowest IP address, as specified by the BGP router ID.
- If the originator or router ID is the same for multiple paths, prefer the path with the minimum cluster list length.
- Prefer the path that comes from the lowest neighbor address.

BGP Multipath

BGP Multipath allows installation into the IP routing table of multiple equal-cost BGP paths to the same destination prefix. Traffic to the destination prefix is then shared across all installed paths.

These paths are installed in the table together with the best path for load-sharing. BGP Multipath does not affect best-path selection. For example, a router still designates one of the paths as the best path, according to the algorithm, and advertises this best path to its BGP peers.

In order to be candidates for multipath, paths to the same destination need to have these characteristics equal to the best-path characteristics:

- Weight
- Local preference
- AS-PATH length
- Origin code
- Multi Exit Discriminator (MED)
- One of these:

- Neighboring AS or sub-AS (before the addition of the BGP Multipaths)
- AS-PATH (after the addition of the BGP Multipaths)

Some BGP Multipath features put additional requirements on multipath candidates:

- The path should be learned from an external or confederation-external neighbor (eBGP).
- The IGP metric to the BGP next hop should be equal to the best-path IGP metric.

These are the additional requirements for internal BGP (iBGP) multipath candidates:

- The path should be learned from an internal neighbor (iBGP).
- The IGP metric to the BGP next hop should be equal to the best-path IGP metric, unless the router is configured for unequal-cost iBGP multipath.

BGP inserts up to *n* most recently received paths from multipath candidates into the IP routing table, where *n* is the number of routes to install to the routing table, as specified when you configure BGP Multipath. The default value, when multipath is disabled, is 1.

For unequal-cost load balancing, you can also use BGP Link Bandwidth.

Note The equivalent next-hop-self is performed on the best path that is selected among eBGP multipaths before it is forwarded to internal peers.

Requirements and Prerequisites for BGP

Model Support

Threat Defense

Firewall Threat Defense Virtual

Supported Domains

Any

User Roles

Admin

Network Admin

Guidelines for BGP

Firewall Mode Guidelines

Does not support transparent firewall mode. BGP is supported only in routed mode.

IPv6 Guidelines

Supports IPv6.

Additional Guidelines

- For BGP, the next hop IP address for the routes is the network IP address and not 0.0.0.0.
- The system does not add route entry for the IP address received over PPPoE in the CP route table. BGP always looks into CP route table for initiating the TCP session, hence BGP does not form TCP session.

  Thus, BGP over PPPoE is not supported.
- BGP is not supported on management-only or BVI interfaces.
- To avoid adjacency flaps due to route updates being dropped if the route update is larger than the minimum MTU on the link, ensure that you configure the same MTU on the interfaces on both sides of the link.
- BGP with PATH MTU (PMTU) can cause adjacency flaps if MTU discovery fails, especially with ECMP routing. Hence, be cautious while using BGP, PMTU, and ECMP as packet drops can occur if MTU discovery fails due to any reason.
- The BGP table of the member unit is not synchronized with the control unit table. Only its routing table is synchronized with the control unit routing table.
- When you configure a route-based site-to-site VPN using static or dynamic VTI interfaces, ensure that the value of the TTL hop is more than one if you use BGP as the routing protocol.

Configure BGP

To configure BGP, see the following topics:

Procedure

Step 1 Configure BGP Basic Settings, on page 1282

Step 2 Configure BGP General Settings, on page 1284

Step 3 Configure BGP Neighbor Settings, on page 1285

Step 4 Configure BGP Aggregate Address Settings, on page 1289

Step 5 Configure BGPv4 Filtering Settings, on page 1290

Note
The Filtering section is applicable only to IPv4 settings

Step 6 Configure BGP Network Settings, on page 1291

Step 7 Configure BGP Redistribution Settings, on page 1291

Step 8 Configure BGP Route Injection Settings, on page 1292

Step 9 Configure BGP Route Import/Export Settings, on page 1293

Configure BGP Basic Settings

You can set many basic settings for BGP.

For a device using virtual routing, the basic settings described in this section must be configured in the **BGP** page under **General Settings**. For more information, see Modifications to the Firewall Management Center Web Interface - Routing Page, on page 1177.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select **Routing**.

Step 3 (For a virtual-router-aware device) Under **General Settings**, click **BGP**.

Step 4 Check the **Enable BGP** check box to enable the BGP routing process.

Step 5 In the **AS Number** field, enter the autonomous system (AS) number for the BGP process. The AS number internally includes multiple autonomous numbers. The AS number can be from 1 to 4294967295 or from 1.0 to 65535.65535. The AS number is a uniquely assigned value, that identifies each network on the Internet.

Step 6 In the **Router ID** drop-down list, choose Automatic or Manual (appears for non-cluster and a cluster in spanned etherchannel mode) or Cluster Pool (appears for a cluster in individual interface mode). If you choose Automatic, the highest-level IP address on the Firewall Threat Defense device is used as the router ID. If you choose Manual, enter the IP address in the **IP Address** field. If you choose Cluster Pool, enter the cluster pool value in the **Cluster Pool** field. For information on creating the cluster pool address, see Address Pools, on page 1374.

Step 7 To use a fixed router ID, choose Manual and enter an IPv4 address in the **IP Address** field. The default value is Automatic. For a virtual router-aware device, you can override the router ID settings in the **Virtual Routers** > **BGP** page.

Step 8 (Optional) Edit the various BGP settings, starting with **General**. The defaults for these settings are appropriate in most cases, but you can adjust them to fit the needs of your network. Click **Edit** (✎) to edit the settings in the group:

a) Enter a **Scanning Interval** for BGP routers for next-hop validation. Valid values are from 5 to 60 seconds. The default value is 60.

b) Enter the **Number of AS numbers in AS_PATH attribute**. An AS _PATH attribute is a sequence of intermediate AS numbers between source and destination routers that form a directed route for packets to travel. Valid values are between 1 and 254. The default value is None.

c) Check the **Log Neighbor Changes** check box to enable logging of BGP neighbor changes (up or down) and resets. This helps in troubleshooting network connectivity problems and measuring network stability. This is enabled by default.

d) Check the **Use TCP Path MTU Discovery** check box to use the Path MTU determining technique to determine the maximum transmission unit (MTU) size on the network path between two IP hosts. This avoids IP fragmentation. This is enabled by default.

e) Check the **Reset session upon Failover** check box to reset the external BGP session immediately upon link failure. This is enabled by default.

f) Check the **Enforce that the first AS is peer's AS for EBGP routes** check box to discard incoming updates received from external BGP peers that do not list their AS number as the first segment in the AS_PATH attribute. This prevents a mis-configured or unauthorized peer from misdirecting traffic by advertising a route as if it was sourced from another autonomous system. This is enabled by default.

g) Check the **Use dot notation for AS number** check box to split the full binary 4-byte AS number into two words of 16 bits each, separated by a dot. AS numbers from 0-65553 are represented as decimal numbers and AS numbers larger than 65535 are represented using the dot notation. This is disabled by default.

h) Click **OK**.

Step 9 (Optional) Edit the **Best Path Selection** section:

a) Enter a value for **Default Local Preference** between 0 and 4294967295. The default value is 100. Higher values indicate higher preference. This preference is sent to all routers and access servers in the local autonomous system.

b) Check the **Allow comparing MED from different neighbors** check box to allow the comparison of Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems. This is disabled by default.

c) Check the **Compare Router ID for identical EBGP paths** check box to compare similar paths received from external BGP peers during the best path selection process and switch the best path to the route with the lowest router ID. This is disabled by default.

d) Check the **Pick the best MED path among paths advertised from the neighboring AS** check box to enable MED comparison among paths learned from confederation peers. The comparison between MEDs is made only if no external autonomous systems are there in the path. This is disabled by default.

e) Check the **Treat missing MED as the least preferred one** check box to consider the missing MED attribute as having a value of infinity, making the path the least desirable; therefore, a path with a missing MED is least preferred. This is disabled by default.

f) Click **OK**.

Step 10 (Optional) Edit the **Neighbor Timers** section:

a) Enter the time interval for which the BGP neighbor remains active after not sending a keepalive message in the **Keep alive interval** field. At the end of this keepalive interval, the BGP peer is declared dead, if no messages are sent. The default value is 60 seconds.

b) Enter the time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured in the **Hold time** field. The default value is 180 seconds. Specify a value from 0 to 65535.

c) (Optional) Enter the minimum time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured in the **Min Hold time** field. Specify a value from 3 to 65535.

Note
A hold time of less than 20 seconds increases the possibility of peer flapping.

d) Click **OK**.

Step 11 In the **Next Hop** section, optionally select the **Enable address tracking** check box to enable BGP next hop address tracking and enter the **Delay Interval** between checks on updated next-hop routes installed in the routing table. Click **OK**.

Note
The **Next Hop** section is applicable only to IPv4 settings.

Step 12 (Optional) Edit the **Graceful Restart** section:

Note
This section is available only when the Firewall Threat Defense device is in failover or spanned cluster mode. This is done so that there is no drop in packets in the traffic flow, when one of the devices in the failover setup fails.

a) Check the **Enable Graceful Restart** checkbox to enable Firewall Threat Defense peers to avoid a routing flap following a switchover.
b) Specify the time duration that Firewall Threat Defense peers will wait to delete stale routes before a BGP open message is received in the **Restart Time** field. The default value is 120 seconds. Valid values are between 1 and 3600 seconds.
c) Enter the time duration that the Firewall Threat Defense will wait before deleting stale routes after an end of record (EOR) message is received from the restarting Firewall Threat Defense in the **Stalepath Time** field. The default value is 360 seconds. Valid values are between 1 and 3600 seconds.
d) Click **OK**.

Step 13 Click **Save**.

Step 14 To view the BGP basic settings, from the virtual routers drop-down, select the desired router, and then click **BGP**.

This page displays the basic settings that are configured in the **Settings** page. You can edit the router ID settings on this page.

Step 15 To edit the router ID settings, modify the IP address in the **IP Address** fields. The modified value overrides the router ID settings that were configured in the **BGP** page under **General Settings**.

Configure BGP General Settings

Configure Route maps, Administrative Route Distances, Synchronization, Next-hop, and packet forwarding. The defaults for these settings are appropriate in most cases, but you can adjust them to fit the needs of your network.

Procedure

Step 1 On the **Device Management** page, click **Routing**.

Step 2 (For a virtual-router-aware device) From the virtual routers drop-down, select the virtual router for which you are configuring BGP.

Step 3 Choose **BGP** > **IPv4** or **IPv6**.

Step 4 Click **General**.

Step 5 In **General**, update the following sections:

a) In the **Settings** section, enter or select a **Route Map** object and click **OK**.

Note
The **Route Map** field is applicable only to IPv4 settings.

b) In the **Administrative Route Distances** section, update the following as required, and click **OK**:

- **External** — Enter the administrative distance for external BGP routes. Routes are external when learned from an external autonomous system. The range of values for this argument are from 1 to 255. The default value is 20.
- **Internal** — Enter administrative distance for internal BGP routes. Routes are internal when learned from peer in the local autonomous system. The range of values for this argument are from 1 to 255. The default value is 200.

- **Local** — Enter administrative distance for local BGP routes. Local routes are those networks listed with a network router show command, often as back doors, for the router or for the networks that is being redistributed from another process. The range of values for this argument are from 1 to 255. The default value is 200.

c) In the **Routes and Synchronization** section, update the following as required, and click **OK**:

- (Optional) **Generate default routes** — Check the check box of this option to configure default-information originate.
- (Optional) **Summarize subnet routes into network-level routes** — Check the check box of this to configure automatic summarization of subnet routes into network-level routes. This check box is applicable only to IPv4 settings.
- (Optional) **Advertise inactive routes** — Check the check box of this to advertise routes that are not installed in the routing information base (RIB).
- (Optional) **Synchronize between BGP and IGP system** — Check the check box of this to enable synchronization between BGP and your Interior Gateway Protocol (IGP) system. Usually, a BGP speaker does not advertise a route to an external neighbor unless that route is local or exists in the IGP. This feature allows routers and access servers within an autonomous system to have the route before BGP makes it available to other autonomous systems.
- (Optional) **Redistribute IBGP into IGP** — Check the check box of this to configure iBGP redistribution into an interior gateway protocol (IGP), such as OSPF.

d) In the **Forward Packets over Multiple Paths** section, update the following as required and click **OK**:

- (Optional) **Number of Paths** — Enter the maximum number of Border Gateway Protocol routes that can be installed in a routing table. The range of values are from 1 to 8. The default value is 1.
- (Optional) **IBGP Number of Paths** — Enter the maximum number of parallel internal Border Gateway Protocol (iBGP) routes that can be installed in a routing table. The range of values are from 1 to 8. The default value is 1.

Step 6 Click **Save**.

Configure BGP Neighbor Settings

A BGP router must connect with each of its peers before exchanging updates. These peers are called BGP neighbors. Perform this procedure to define BGP IPv4 or IPv6 neighbors and neighbor settings.

Procedure

Step 1 On the Device Management page, click **Routing**.

Step 2 (For a virtual router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring BGP.

Step 3 Choose **BGP** > **IPv4** or **IPv6**.

Step 4 Click **Neighbor**.

Step 5 Click **Add** to define BGP neighbors and neighbor settings.

Step 6 Enter the BGP neighbor **IP address**. This IP address is added to the BGP neighbor table. When you are configuring BGP IPv6 on static VTI, enter the virtual tunnel IP address of the neighbor.

Step 7 Choose the BGP neighbor **Interface**.

Note
The **Interface** field is only applicable to IPv6 settings.

Step 8 Enter the autonomous system to which the BGP neighbor belongs, in the **Remote AS** field.

Step 9 Check the **Enabled address** check box to enable communication with this BGP neighbor. Further neighbor settings will be configured only if the Enabled address check box is selected.

Step 10 (Use this option only when the routing does not result in a loop.) Check the **AS Override** check box to enable overriding of the AS number of the originating router with the AS number of the sending BGP router.

Note
Loop prevention in BGP is achieved by verifying the AS number in the AS Path. BGP rejects route updates when the AS Path attribute contain its own AS number. However, in circumstances where the same ASN is used by discontiguous network segments, to establish end-to-end reacheability, you can opt to override the AS number with that of the sending BGP router.

Step 11 (Optional) Check the **Shutdown administratively** check box to disable a neighbor or peer group.

Step 12 (Optional) Check the **Configure graceful restart (failover / spanned mode)** check box to enable configuration of the BGP graceful restart capability for this neighbor. After selecting this option, you must check the **Enable graceful restart** check box to specify whether graceful restart should be enabled or disabled for this neighbor.

Note
- The graceful restart is enabled only when the device is in HA mode or when L2 cluster (all nodes from the same network) is configured.
- The graceful restart option for BGPv6 is enabled only on Firewall Threat Defense Version 7.3+.
- If you configure graceful restart only at General Settings and not at BGP IPv6, the global General Settings configuration persist.
- If you configure graceful restart at General Settings and also at BGP IPv6, the global General Settings configuration is overridden by the BGP IPv6 configuration settings.

Step 13 (Optional) To enable configuration of the BFD support for BGP, from the **BFD Fallover** drop-down list, choose the BFD type—single-hop, multi-hop, or auto-detect-hop. This selection registers the BGP neighbor to receive forwarding path detection failure messages from BFD. Choose None if you do not want to have BFD support.

Step 14 (Optional) Enter a **Description** for the BGP neighbor.

Step 15 (Optional) From the **Update Source**drop-down list, choose an interface to source the BGP packets.

You can choose a loopback address as this interface to overcome path failures. You can also choose any physical, port-channel, or a sub-interface.

Step 16 (Optional) In **Filtering Routes**, use access lists, route maps, prefix lists and AS path filters as required, to distribute BGP Neighbor information. Update the following sections:

a) Choose or select the appropriate incoming or outgoing **Access List** to distribute BGP neighbor information.

Note

Access lists are only applicable to IPv4 settings.

b) Choose or select the appropriate incoming or outgoing **Route Maps** to apply a route map to incoming or outgoing routes.
c) Choose or select the appropriate incoming or outgoing **Prefix List** to distribute BGP neighbor information.
d) Choose or select the appropriate incoming or outgoing **AS path filter** to distribute BGP neighbor information.
e) Check the check box of **Limit the number of prefixes allowed from the neighbor** to control the number of prefixes that can be received from a neighbor.
   - Enter the maximum number of prefixes allowed from a specific neighbor in the **Maximum Prefixes** field.
   - Enter the percentage (of maximum) at which the router starts to generate a warning message in the **Threshold Level** field. Valid values are integers between 1 and 100. The default value is 75.
f) Check the **Control prefixes received from the peer** check box to specify additional controls for the prefixes received from a peer. Do one of the following
   - Check the **Terminate peering when prefix limit is exceeded** check box to stop the BGP neighbor when the prefix limit is reached. Specify the interval after which the BGP neighbor will restart in the **Restart interval** field.
   - Check the **Give only warning message when prefix limit is exceeded** check box to generate a log message when the maximum prefix limit is exceeded. Here, the BGP neighbor will not be terminated.
g) Click **OK**.

Step 17 (Optional) In **Routes**, specify miscellaneous Neighbor route parameter. Proceed to update the following:

a) Enter the minimum interval (in seconds) between the sending of BGP routing updates in the **Advertisement Interval** field. Valid values are between 1 and 600.
b) Check the **Remove private AS numbers from outbound routing updates** check box to exclude the private AS numbers from being advertised on outbound routes.
c) Check the **Generate default routes** check box to allow the local router to send the default route 0.0.0.0 to a neighbor to use as a default route. Enter or Select the route map that allows the route 0.0.0.0 to be injected conditionally in the **Route map** field.
d) To add conditionally advertised routes, click Add Row +. In the Add Advertised Route dialog box, do the following:
   1. Add or choose a route map in the **Advertise Map** field, that will be advertised if the conditions of the exist map or the non-exist map are met.
   2. Click **Exist Map** and choose a route map from the Route Map Object Selector. This route map is compared with the routes in the BGP table, to determine whether the advertise map route is advertised.
   3. Click **Non-Exist Map** and choose a route map from the Route Map Object Selector. This route map is compared with the routes in the BGP table, to determine whether the advertise map route is advertised.
   4. Click **OK**.

Step 18 In **Timers**, check the **Set timers for the BGP peer** check box to set the keepalive frequency, hold time and minimum hold time

- **Keep alive interval**—Enter the frequency (in seconds) with which Firewall Threat Defense sends keepalive messages to the neighbor. Valid values are between 0 and 65535. The default value is 60 seconds.
- **Hold time**—Enter the interval (in seconds) after not receiving a keepalive message that Firewall Threat Defense declares a peer dead. Valid values are between 0 and 65535. The default value is 180 seconds.
- **Min hold time**—(Optional) Enter the minimum interval (in seconds) after not receiving a keepalive message that Firewall Threat Defense declares a peer dead. Valid values are between 3 and 65535. The default value is 3 seconds.

Note
A hold time of less than 20 seconds increases the possibility of peer flapping.

Step 19 In **Advanced**, update the following:

a) (Optional) Check the **Enable Authentication** check box to enable MD5 authentication on a TCP connection between two BGP peers.

1. Choose an encryption type from the **Enable Encryption** drop-down list.
2. Enter a password in the **Password** field. Reenter the password in the **Confirm Password** field. The password is case-sensitive and can be up to 25 characters long when the service password-encryption command is enabled and up to 81 characters long when the service password-encryption command is not enabled. The string can contain any alphanumeric characters, including spaces.

Note
You cannot specify a password in the format number-space-anything. The space after the number can cause authentication to fail.

b) (Optional) Select the **Send Community attribute to this neighbor** check box to specify that communities attributes should be sent to the BGP neighbor

c) (Optional) Select the **Use FTD as next hop for this neighbor** check box to configure the router as the next-hop for a BGP speaking neighbor or peer group.

d) Select the **Disable Connection Verification** check box to disable the connection verification process for eBGP peering sessions that are reachable by a single hop but are configured on a loopback interface or otherwise configured with a non-directly connected IP address. When deselected (default), a BGP routing process will verify the connection of single-hop eBGP peering session (TTL=254) to determine if the eBGP peer is directly connected to the same network segment by default. If the peer is not directly connected to same network segment, connection verification will prevent the peering session from being established.

e) Select **Allow connections with neighbor that is not directly connected** to accept and attempt BGP connections to external peers residing on networks that are not directly connected. (Optional) Enter the time-to-live in the **TTL hops**field. Valid values are between 1 and 255. Alternately, select **Limited number of TTL hops to neighbor**, to secure a BGP peering session. Enter the maximum number of hops that separate eBGP peers in the **TTL hops** field. Valid values are between 1 and 254.

f) (Optional) Select the **Use TCP MTU path discovery** check box to enable a TCP transport session for a BGP session.

g) Choose the TCP connection mode from the **TCP Transport Mode** drop-down list. Options are Default, Active, or Passive.

h) (Optional) Enter a **Weight** for the BGP neighbor connection.

i) Select the **BGP Version** that Firewall Threat Defense will accept from the drop-down list. The version can be set to 4-Only to force the software to use only Version 4 with the specified neighbor. The default is to use Version 4 and dynamically negotiate down to Version 2 if requested.

Step 20 Update **Migration**, only if AS migration is considered.

Note
The AS migration customization should be removed after transition has been completed.

a) (Optional) Check the **Customize the AS number for routes received from the neighbor** check box to customize the AS_PATH attribute for routes received from an eBGP neighbor.
b) Enter the local autonomous system number in the **Local AS number** field. Valid values are any valid autonomous system number from 1 to 4294967295 or 1.0 to 65535.65535.
c) (Optional) Check the **Do not prepend local AS number to routes received from neighbor** check box to prevent the local AS number from being prepended to any routes received from eBGP peer.
d) (Optional) Check the **Replace real AS number with local AS number in routes received from neighbor** check box to replace the real autonomous system number with the local autonomous system number in the eBGP updates. The autonomous system number from the local BGP routing process is not prepended.
e) (Optional) Check the **Accept either real AS number or local AS number in routesreceived from neighbor** check box to configure the eBGP neighbor to establish a peering session using the real autonomous system number (from the local BGP routing process) or by using the local autonomous system number.

Step 21 Click **OK**.

Step 22 Click **Save**.

Configure BGP Aggregate Address Settings

BGP neighbors store and exchange routing information and the amount of routing information increases as more BGP speakers are configured. Route aggregation is the process of combining the attributes of several different routes so that only a single route is advertised. Aggregate prefixes use the classless interdomain routing (CIDR) principle to combine contiguous networks into one classless set of IP addresses that can be summarized in routing tables. As a result fewer routes need to be advertised. Use the Add/Edit Aggregate Address dialog box to define the aggregation of specific routes into one route.

Procedure

Step 1 When editing the Firewall Threat Defense device, click **Routing**.

Step 2 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.

Step 3 Choose **BGP** > **IPv4** or **IPv6**.

Step 4 Click **Add Aggregate Address**.

Step 5 Enter a value for the aggregate timer (in seconds) in the **Aggregate Timer** field. Valid values are 0 or any value between 6 and 60. The default value is 30.

Step 6 Click (+)**Add** and update the **Add Aggregate Address** dialog box:

a) **Network** — Enter an IPv4 address or select the desired network/hosts objects.

b) **Attribute Map** — (Optional) Enter or select the route map used to set the attribute of the aggregate route.
c) **Advertise Map** — (Optional) Enter or select the route map used to select the routes to create AS_SET origin communities.
d) **Suppress Map** — (Optional) Enter or select the route map used to select the routes to be suppressed.
e) **Generate AS set path Information** — (Optional) Check the check box to enable generation of autonomous system set path information.
f) **Filter all routes from updates** — (Optional) Check the check box to filter all more-specific routes from updates.
g) Click **OK**.

What to do next

- For BGPv4 settings, proceed to Configure BGPv4 Filtering Settings, on page 1290.
- For BGPv6 settings, proceed to Configure BGP Network Settings, on page 1291.

Configure BGPv4 Filtering Settings

Filtering settings are used to filter routes or networks received in incoming BGP updates. Filtering is used to restrict routing information that the router learns or advertises.

Before you begin

Filtering is only applicable for a BGP IPv4 routing policy.

Procedure

Step 1 On the Device Management page, click **Routing**.

Step 2 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.

Step 3 Choose **BGP** > **IPv4**.

Step 4 Click **Filtering**.

Note
The **Filtering** field is applicable only to IPV4 settings.

Step 5 Click (+)**Add** and update the **Add Filter** dialog box:

a) **Access List**— Choose an access control list that defines which networks are to be received and which are to be suppressed in routing updates.
b) **Direction**— (Optional) Choose a direction that specifies if the filter should be applied to inbound updates or outbound updates.
c) **Protocol**— (Optional) Choose the routing process for which you want to filter: None, BGP, Connected, OSPF, RIP, or Static.
d) **Process ID**— (Optional) Enter the process ID for the OSPF routing protocol.
e) Click **OK**.

Step 6 Click **Save**.

Configure BGP Network Settings

Network settings are used to add networks that will be advertised by the BGP routing process and route maps that will be examined to filter the networks to be advertised.

Procedure

Step 1 On the **Device Management** page, click **Routing**.

Step 2 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.

Step 3 Choose **BGP** > **IPv4** or **IPv6**.

Step 4 Click **Networks**.

Step 5 Click **Add** and update the **Add Networks** dialog box:

a) **Network**— Choose the network to be advertised by the BGP routing processes.

Note
For a network prefix to be advertised, a route to the device must exist on the routing table.

To add a new network object, see Creating Network Objects, on page 1403.

b) (Optional) **Route Map**— Enter or choose a route map that should be examined to filter the networks to be advertised. If not specified, all networks are redistributed. To add a new route map object, see Route Map, on page 1426.

c) Click **OK**.

Step 6 Click **Save**.

Configure BGP Redistribution Settings

Redistribution settings allow you to define the conditions for redistributing routes from another routing domain into BGP.

Procedure

Step 1 On the **Device Management** page, click **Routing**.

Step 2 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.

Step 3 Choose **BGP** > **IPv4** or **IPv6**.

Step 4 Click **Redistribution**.

Step 5 Click **Add** and update the **Add Redistribution** dialog:

a) **Source Protocol**— Select the protocol from which you want to redistribute routes into the BGP domain from the Source Protocol drop-down list.

Note
User-defined virtual routers does not support redistributing traffic from RIP.

b) **Process ID**— Enter the identifier for the selected source protocol. Applies to the OSPF protocol. For devices using virtual routing, this drop-down lists the process ID assigned for the virtual router for which you are configuring the BGP settings.
c) **Metric**— (Optional) Enter a metric for the redistributed route.
d) **Route Map**— Enter or select a route map that should be examined to filter the networks to be redistributed.

If not specified, all networks are redistributed. To create a new route map object, click **Add** (+). See Configure Route Map Entry for the procedure to add a new route map.
e) **Match**— The conditions used for redistributing routes from one routing protocol to another. The routes must match the selected condition to be redistributed. You can choose one or more of the following match conditions. These options are enabled only when OSPF is chosen as the Source Protocol.

- Internal
- External 1
- External 2
- NSSA External 1
- NSSA External 2

f) Click **OK**.

Configure BGP Route Injection Settings

Route injection settings allow you to define the routes to be conditionally injected into the BGP routing table.

Procedure

Step 1 On the **Device Management** page, click **Routing**.
Step 2 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.
Step 3 Choose **BGP** > **IPv4** or **IPv6**.
Step 4 Click **Route Injection**.
Step 5 Click **Add** and update the **Add Route Injection** dialog box:

a) **Inject Map**— Enter or select the route map that specifies the prefixes to inject into the local BGP routing table. To create a new route map object, click **Add** (+). For the procedure to add a new route map, see Configure Route Map Entry.
b) **Exist Map**— Enter or select the route map containing the prefixes that the BGP speaker will track.
c) **Injected routes will inherit the attributes of the aggregate route**— Check this box to configure the injected route to inherit attributes of the aggregate route.

d) Click **OK**.

Step 6 Click **Save**.

Configure BGP Route Import/Export Settings

In BGP, you can implement an inter-virtual-router route leak by importing or exporting routes using the route target extended community of the destination and source virtual routers respectively. You can use a route map to filter the desired route targets instead of leaking the entire routing table. You can also leak the routes of global virtual router to user-defined virtual routers and vice versa.

- You can configure BGP to leak routes between two user-defined virtual routers using the route target extended communities:
  - Tag the routes with the route targets from the source virtual router using route target export.
  - Import the routes that are matching the route targets in to the destination virtual router using route target import.
  - Optionally, you can filter routes from source virtual router or to destination virtual router using export or import route maps respectively. You can configure route map with match extended community list for filtering the routes. Similarly, you can configure route map with set extended community route targets to tag the routes with the route target extended community.
- To import routes from the global virtual router to a user-defined virtual router, specify the IPv4/IPv6 route map in Global Virtual Router Import Route Map to import to the user-defined virtual router.
- To export routes from a user-defined virtual router to the global virtual router, in addition to exporting the route targets, you can also specify the Global Virtual Router Export Route Map to export from the user-defined virtual router.

The BGP inter-virtual-router route leaking supports both ipv4 and ipv6 prefixes.

Before you begin

- Create virtual routers.
- Enable BGP on the virtual router.
- Configure BGP, on page 1281.

Procedure

Step 1 On the Device Management page, click **Routing**.

Step 2 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.

Step 3 Choose **BGP** > **IPv4** or **IPv6**.

Step 4 (Supported only for only virtual routers) Click **Route Import/Export**.

Step 5 In the **Route Targets Import** field, enter the route target extended community that you want to match for the routes to be imported. On deployment, the routes of the destination virtual router that matches this value is imported to the source virtual router's BGP table.

Note
- The route target must be in **ASN:nn** format.
- You can enter multiple route targets as comma separated values.
- This value can range from 0:1 to 65534:65535.

Step 6 In the **Route Targets Export** field, enter the route target extended community to tag the source virtual router's routes with the route target value. On deployment, the routes of the source virtual router are tagged with this value.

Note
- The route target must be in **ASN:nn** format.
- You can enter multiple route targets as comma separated values.
- This value can range from 0:1 to 65534:65535.

Step 7 Route maps help you to narrow down the routes to be shared instead of leaking the entire routing table. Route map filtering is applied on the list of routes that are obtained with the specified route target values:

a) (Optional) Under **User Virtual Router**, choose the route map from the **Import Route Map** drop-down list to filter the routes at the destination virtual router.

Note
The user virtual router import route map is effective only when the route targets import is configured.

b) (Optional) Under **User Virtual Router**, choose the route map from the **Export Route Map** drop-down list to filter the routes at the source virtual router before the routes are exported to other virtual routers.

Note
You can use the match and set clauses in the route map with the route target extended community lists for filtering based on other criteria or tagging the routes with the route target community values. For more information, see Route Map, on page 1426.

Step 8 To share the routes between a user-defined virtual router and global virtual router, specify the route map under the **Global Virtual Router**:

a) To leak the global virtual router routes to the user-defined virtual router, select the route map from the **Import Route Map** drop-down list. The IPv4 or IPv6 route map is imported to the user-defined virtual router.

b) To leak the user-defined virtual router routes to the global virtual router, select the route map from the **Export Route Map** drop-down list. The IPv4 or IPv6 route map is exported to the global virtual router.

Note
You must specify the route targets for export apart from specifying the route map.

Note
You can use the match clause of the route map object to filter the routes for leaking. For more information, see Route Map, on page 1426.

Step 9 Follow the procedure (Step 3 to Step 8) to configure relevant BGP route import and export settings for other virtual routers as well.

Step 10 Click **Save** and **Deploy**.

When the packets flow into the ingress virtual router, BGP imports the routes from the destination virtual routers that have the matching route target value and if a route map is also configured, the routes are further filtered and used to identify the best path routes for routing the packets.

History for BGP

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
BGP AS Override	7.7	7.7	You can enable the receiving BGP router to override the AS number of the originating router (that is the same as the receiving router) with the AS number of the sending BGP router if the same AS number is being used by discontiguous network segments. New/modified screens: **Routing** > **BGP** > **IPv4 or IPv6** > **Add/Edit Neighbor**.
Graceful restart support on BGPv6	7.4	Any	You can configure graceful restart on BGPv6 for Secure Firewall Threat Defense version 7.3 and later. New/modified screens: **Routing** > **BGP** > **IPv6** > **Add/Edit Neighbor** .
Loopback interface support for BGP	7.4	Any	You can use a loopback interface for BGP. New/modified screens: **Routing** > **BGP** > **IPv4 or IPv6** > **Add/Edit Neighbor** .
BGP configuration to interconnect virtual routers	7.1	Any	You can configure BGP settings to dynamically leak routes among user-defined virtual routers, and between global virtual router and user-defined virtual routers. The import and export routes feature was introduced to exchange routes among the virtual routers by tagging them with route targets and optionally, filtering the matched routes with route maps. This BGP feature is accessible only when you select a user-defined virtual router. New/modified screens: For a selected user-defined virtual router, **Devices** > **Device Management** > **Routing** > **BGPv4/v6** > **Route Import/Export** tab.
BGPv6 support for user-defined virtual routers	7.1	Any	Secure Firewall Threat Defense now supports configuring BGPv6 on user-defined virtual routers. New/modified screens: For a selected user-defined virtual router, **Devices** > **Device Management** > **Routing** > **BGPv6** page.

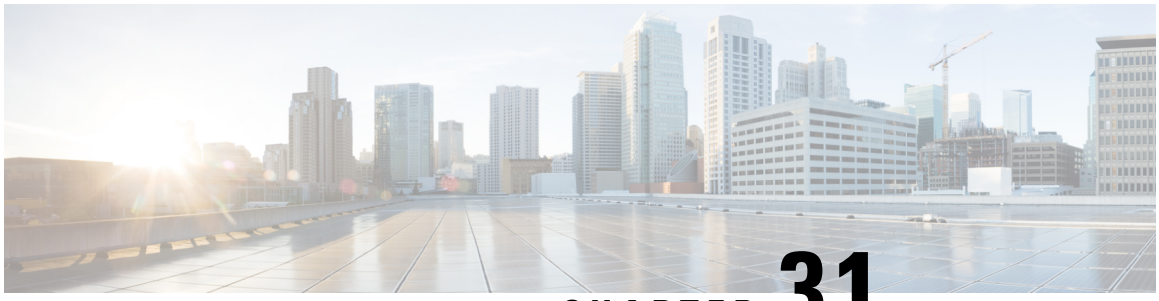

CHAPTER 31

RIP

This chapter describes how to configure the Firewall Threat Defense to route data, perform authentication, and redistribute routing information, using the Routing Information Protocol (RIP). For a device using virtual routing, you can configure RIP only for its global virtual router and not for its user-defined virtual router.

- About RIP, on page 1297
- Requirements and Prerequisites for RIP, on page 1299
- Guidelines for RIP, on page 1299
- Configure RIP, on page 1300

About RIP

The Routing Information Protocol, or RIP, as it is more commonly called, is one of the most enduring of all routing protocols. RIP has four basic components: routing update process, RIP routing metrics, routing stability, and routing timers. Devices that support RIP send routing-update messages at regular intervals and when the network topology changes. These RIP packets include information about the networks that the devices can reach, as well as the number of routers or gateways that a packet must travel through to reach the destination address. RIP generates more traffic than OSPF, but is easier to configure.

RIP is a distance-vector routing protocol that uses hop count as the metric for path selection. When RIP is enabled on an interface, the interface exchanges RIP broadcasts with neighboring devices to dynamically learn about and advertise routes.

The Secure Firewall Threat Defense device supports both RIP Version 1 and RIP Version 2. RIP Version 1 does not send the subnet mask with the routing update. RIP Version 2 sends the subnet mask with the routing update and supports variable-length subnet masks. Additionally, RIP Version 2 supports neighbor authentication when routing updates are exchanged. This authentication ensures that the Secure Firewall Threat Defense device receives reliable routing information from a trusted source.

RIP has advantages over static routes because the initial configuration is simple, and you do not need to update the configuration when the topology changes. The disadvantage to RIP is that there is more network and processing overhead than in static routing.

Routing Update Process

RIP sends routing-update messages at regular intervals and when the network topology changes. When a router receives a routing update that includes changes to an entry, it updates its routing table to reflect the new route. The metric value for the path is increased by 1, and the sender is indicated as the next hop. RIP

routers maintain only the best route (the route with the lowest metric value) to a destination. After updating its routing table, the router immediately begins transmitting routing updates to inform other network routers of the change. These updates are sent independently of the regularly scheduled updates that RIP routers send.

RIP Routing Metric

RIP uses a single routing metric (hop count) to measure the distance between the source and a destination network. Each hop in a path from source to destination is assigned a hop count value, which is typically 1. When a router receives a routing update that contains a new or changed destination network entry, the router adds 1 to the metric value indicated in the update and enters the network in the routing table. The IP address of the sender is used as the next hop.

RIP Stability Features

RIP prevents routing loops from continuing indefinitely by implementing a limit on the number of hops allowed in a path from the source to a destination. The maximum number of hops in a path is 15. If a router receives a routing update that contains a new or changed entry, and if increasing the metric value by 1 causes the metric to be infinity (that is, 16), the network destination is considered unreachable. The downside of this stability feature is that it limits the maximum diameter of a RIP network to less than 16 hops.

RIP includes a number of other stability features that are common to many routing protocols. These features are designed to provide stability despite potentially rapid changes in network topology. For example, RIP implements the split horizon and hold-down mechanisms to prevent incorrect routing information from being propagated.

RIP Timers

RIP uses numerous timers to regulate its performance. Following are the timer stages for RIP:

- Update—The routing-update timer is the interval between periodic routing updates. This is how often the device sends routing updates. Generally, it is set to 30 seconds, with a small random amount of time added whenever the timer is reset. This is done to help prevent congestion, which could result from all routers simultaneously attempting to update their neighbors.
- Invalid—Each routing table entry has a route-timeout timer associated with it. This is the number of seconds since the device received the last valid update. When the route-timeout timer expires, the route is marked invalid but is retained in the table until the route-flush timer expires. Once this timer expires, the route goes into holddown. The default is 180 seconds (3 minutes).
- Holddown—The holddown period is the number of seconds the system waits before accepting any new updates for the route that is in holddown (that is, routes that have been marked invalid). The default is 180 seconds (3 minutes).
- Flush—The route-flush timer is the number of seconds since the system received the last valid update until the route is discarded and removed from the routing table. The default is 240 seconds (4 minutes).

As an example, when the interface on an adjacent router goes down, the system no longer receives routing updates from the adjacent router. At this time, the Invalid and Flush timers start increasing. In the first 180 seconds, nothing will happen. After 180 seconds, the invalid timer expires, making the route invalid, and the Holddown timer starts and holds the route for another 60 seconds. If there is still no update regarding the interface status on the adjacent router (that is, it is still down), then the route enters into the Flush state where in total the system has waited for 240 seconds from the last update (180 seconds for the Invalid timer and 60

seconds for Holddown timer), and the system flushes the route. Even if the adjacent routers interface comes up immediately, the system does not accept a routing update until the Holddown timer completes the remaining 120 seconds.

Requirements and Prerequisites for RIP

Model Support

Threat Defense

Firewall Threat Defense Virtual

Supported Domains

Any

User Roles

Admin

Network Admin

Guidelines for RIP

IPv6 Guidelines

Does not support IPv6.

Additional Guidelines

The following information applies to RIP Version 2 only:

- If using neighbor authentication, the authentication key and key ID must be the same on all neighbor devices that provide RIP Version 2 updates to the interface.
- With RIP Version 2, the Secure Firewall Threat Defense device transmits and receives default route updates using the multicast address 224.0.0.9. In passive mode, it receives route updates at that address.
- When RIP Version 2 is configured on an interface, the multicast address 224.0.0.9 is registered on that interface. When a RIP Version 2 configuration is removed from an interface, that multicast address is unregistered.

Limitations

- The Secure Firewall Threat Defense device cannot pass RIP updates between interfaces.
- RIP Version 1 does not support variable-length subnet masks.
- RIP has a maximum hop count of 15. A route with a hop count greater than 15 is considered unreachable.
- RIP convergence is relatively slow compared to other routing protocols.

- You can only enable a single RIP process on the Secure Firewall Threat Defense device.

Configure RIP

RIP is a distance-vector routing protocol that uses hop count as the metric for path selection.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select **Routing**.

Step 3 Select **RIP** from the table of contents.

Step 4 Check the **Enable RIP** check box to configure the RIP settings.

Step 5 Choose the RIP versions for sending and receiving RIP updates from the **RIP Version** drop-down list.

Step 6 (Optional) Check the **Generate Default Route** check box to generate a default route for distribution, based on the route map that you specify.

a) Specify a route map name to use for generating default routes, in the **Route Map** field.
The default route 0.0.0.0/0 is generated for distribution over a certain interface , when the route map, specified in the **Route Map** field, is present.

Step 7 When Send and Receive Version 2 is the chosen RIP Version, the **Enable Auto Summary** option is available. When the **Enable Auto Summary** check box is checked, automatic route summarization is enabled. Disable automatic summarization if you must perform routing between disconnected subnets. When automatic summarization is disabled, subnets are advertised.

Note
RIP Version 1 always uses automatic summarization—you cannot disable it.

Step 8 Click **Networks**. Define one or more networks for RIP routing. Enter IP address(es), or enter or select the desired Network/Hosts objects. There is no limit to the number of networks you can add to the security appliance configuration. Any interface that belongs to a network defined by this command, will participate in the RIP routing process. The RIP routing updates will be sent and received only through interfaces on the specified networks. Also, if the network of an interface is not specified, the interface will not be advertised in any RIP updates.

Note
RIP only supports IPv4 objects.

Step 9 (Optional) Click **Passive Interface**. Use this option to specify passive interfaces on the appliance, and by extension the active interfaces. The device listens for RIP routing broadcasts on passive interfaces, using that information to populate its routing tables, but does not broadcast routing updates on passive interfaces. Interfaces that are not designated as passive, receive and send updates.

Step 10 Click **Redistribution** to manage redistribution routes. These are the routes that are being redistributed from other routing processes into the RIP routing process.

a) Click **Add** to specify redistribution routes.

b) Choose the routing protocol to redistribute into the RIP routing process, in the **Protocol** drop-down list.

Note

For the OSPF protocol, specify a process ID. Similarly, specify an AS path for BGP. When you choose the Connected option in the **Protocol** drop-down list, you can redistribute, directly connected networks into the RIP routing process.

c) (Optional) If you are redistributing OSPF routes into the RIP routing process, you can select specific types of OSPF routes to redistribute in the **Match** drop-down list . Ctrl-click to select multiple types:

- **Internal** – Routes internal to the autonomous system (AS) are redistributed.
- **External 1** – Type 1 routes external to the AS are redistributed.
- **External 2** – Type 2 routes external to the AS are redistributed.
- **NSSA External 1** – Type 1 routes external to a not-so-stubby area (NSSA) are redistributed.
- **NSSA External 2** – Type 2 routes external to an NSSA are redistributed

Note
The default is match Internal, External 1, and External 2

d) Select the RIP metric type to apply to the redistributed routes in the **Metric** drop-down list. The two choices are:

- **Transparent** – Use the current route metric
- **Specified Value** – Assign a specific metric value. Enter a specific value from 0-16, in the **Metric Value** field.
- **None** – No metric is specified. Do not use any metric value, to apply to redistributed routes.

Note
None option is applicable only for Static and Connected protocols.

e) (Optional) Enter the name of a route map that must be satisfied, in the **Route Map** field before the route can be redistributed into the RIP routing process. Routes are redistributed only if IP address matches an allow statement in the route map address list. To create a new route map object, click **Add** (+). See Configure Route Map Entry for the procedure to add a new route map.

f) Click **OK**.

Step 11 (Optional) Click **Filtering** to manage filters for the RIP policy. In this section, filters are used to prevent routing updates through an interface, control the advertising of routes in routing updates, control the processing of routing updates and filtering sources of routing updates.

a) Click **Add** to add RIP filters.

b) Select the type of traffic to be filtered - Inbound or Outbound in the **Traffic Direction** field.

Note
If traffic direction is inbound, you can only define an Interface filter.

c) Specify whether the filter is based on an Interface or a Route, by selecting appropriate in the **Filter On** field. If you click **Interface**, enter or choose the name of the interface on which routing updates are to be filtered. If you click **Route**, choose the route type:

- **Static** – Only static routes are filtered.
- **Connected** – Only connected routes are filtered.

- **OSPF** – Only OSPFv2 routes discovered by the specified OSPF process are filtered. Enter the Process ID of the OSPF process to be filtered.
- **BGP** – Only BGPv4 routes discovered by the specified BGP process are filtered. Enter the AS path of the BGP process to be filtered.

d) In the **Access List** field, enter or choose the name of one or more access control lists (ACLs) that define the networks to be allowed or removed from RIP route advertisements. To add a new standard access list object, click **Add** (+) and see Configure Standard ACL Objects, on page 1374.

e) Click **OK**.

Step 12 (Optional) Click **Broadcast** to add or edit interface configurations. Using Broadcast, you can override the global RIP versions to send or receive per interface. You can also define the authentication parameters per interface if you want to implement authentication to ensure valid RIP updates.

a) Click **Add** to add interface configurations.

b) Enter or choose an interface defined on this appliance in the **Interface** field.

c) In the Send option, select the appropriate boxes to specify sending updates using the RIP **Version 1**, **Version 2**, or both. These options let you override, for the specified interface, the global Send versions specified .

d) In the Receive option, select the appropriate boxes to specify accepting updates using the RIP **Version 1**, **Version 2**, or both. These options let you override, for the specified interface, the global Receive versions specified .

e) Select the **Authentication** used on this interface for RIP broadcasts.

- **None** – No authentication
- **MD5** – Employ MD5
- **Clear Text** – Employ clear-text authentication

If you choose MD5 or Clear Text, you must also provide the following authentication parameters.

- **Key ID** – The ID of the authentication key. Valid values are from 0 to 255.
- **Key** – The key used by the chosen authentication method. Can contain up to 16 characters
- **Confirm** – Enter the authentication key again, to confirm

f) Click **OK**.

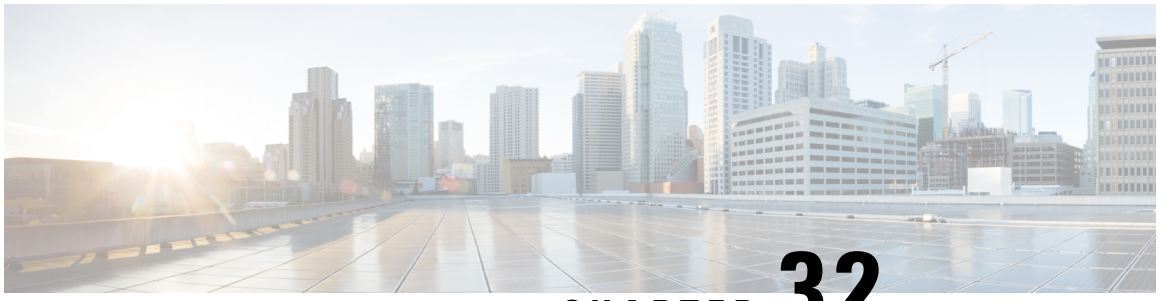

CHAPTER 32

Multicast

This chapter describes how to configure the Secure Firewall Threat Defense device to use the multicast routing protocol.

- About Multicast Routing, on page 1303
- Requirements and Prerequisites for Multicast Routing, on page 1307
- Guidelines for Multicast Routing, on page 1307
- Configure IGMP Features, on page 1308
- Configure PIM Features, on page 1313
- Configure Multicast Routes, on page 1319
- Configure Multicast Boundary Filters, on page 1320

About Multicast Routing

Multicast routing is a bandwidth-conserving technology that reduces traffic by simultaneously delivering a single stream of information to thousands of corporate recipients and homes. Applications that take advantage of multicast routing include videoconferencing, corporate communications, distance learning, and distribution of software, stock quotes, and news.

Multicast routing protocols deliver source traffic to multiple receivers without adding any additional burden on the source or the receivers while using the least network bandwidth of any competing technology. Multicast packets are replicated in the network by Firewall Threat Defense device enabled with Protocol Independent Multicast (PIM) and other supporting multicast protocols, which results in the most efficient delivery of data to multiple receivers possible.

The Firewall Threat Defense device supports both stub multicast routing and PIM multicast routing. However, you cannot configure both concurrently on a single Firewall Threat Defense device.

Note The UDP and non-UDP transports are both supported for multicast routing. However, the non-UDP transport has no FastPath optimization.

IGMP Protocol

IP hosts use the Internet Group Management Protocol (IGMP) to report their group memberships to directly-connected multicast routers. IGMP is used to dynamically register individual hosts in a multicast

group on a particular LAN. Hosts identify group memberships by sending IGMP messages to their local multicast router. Under IGMP, routers listen to IGMP messages and periodically send out queries to discover which groups are active or inactive on a particular subnet.

IGMP uses group addresses (Class D IP address) as group identifiers. Host group address can be in the range of 224.0.0.0 to 239.255.255.255. The address 224.0.0.0 is never assigned to any group. The address 224.0.0.1 is assigned to all systems on a subnet. The address 224.0.0.2 is assigned to all routers on a subnet.

Note When you enable multicast routing on the Firewall Threat Defense device, IGMP Version 2 is automatically enabled on all interfaces.

Query Messages to Multicast Groups

The Firewall Threat Defense device sends query messages to discover which multicast groups have members on the networks attached to the interfaces. Members respond with IGMP report messages indicating that they want to receive multicast packets for specific groups. Query messages are addressed to the all-systems multicast group, which has an address of 224.0.0.1, with a time-to-live value of 1.

These messages are sent periodically to refresh the membership information stored on the Firewall Threat Defense device. If the Firewall Threat Defense device discovers that there are no local members of a multicast group still attached to an interface, it stops forwarding multicast packets for that group to the attached network, and it sends a prune message back to the source of the packets.

By default, the PIM designated router on the subnet is responsible for sending the query messages. By default, they are sent once every 125 seconds.

When changing the query response time, by default, the maximum query response time advertised in IGMP queries is 10 seconds. If the Firewall Threat Defense device does not receive a response to a host query within this amount of time, it deletes the group.

Stub Multicast Routing

Stub multicast routing provides dynamic host registration and facilitates multicast routing. When configured for stub multicast routing, the Firewall Threat Defense device acts as an IGMP proxy agent. Instead of fully participating in multicast routing, the Firewall Threat Defense device forwards IGMP messages to an upstream multicast router, which sets up delivery of the multicast data. When configured for stub multicast routing, the Firewall Threat Defense device cannot be configured for PIM sparse or bidirectional mode. You must enable PIM on the interfaces participating in IGMP stub multicast routing.

The Firewall Threat Defense device supports both PIM-SM and bidirectional PIM. PIM-SM is a multicast routing protocol that uses the underlying unicast routing information base or a separate multicast-capable routing information base. It builds unidirectional shared trees rooted at a single Rendezvous Point (RP) per multicast group and optionally creates shortest-path trees per multicast source.

PIM Multicast Routing

Bidirectional PIM is a variant of PIM-SM that builds bidirectional shared trees connecting multicast sources and receivers. Bidirectional trees are built using a Designated Forwarder (DF) election process operating on each link of the multicast topology. With the assistance of the DF, multicast data is forwarded from sources to the Rendezvous Point (RP), and therefore along the shared tree to receivers, without requiring source-specific state. The DF election takes place during RP discovery and provides a default route to the RP.

Note If the Firewall Threat Defense device is the PIM RP, use the untranslated outside address of the Firewall Threat Defense device as the RP address.

PIM Source Specific Multicast Support

The Firewall Threat Defense device does not support PIM Source Specific Multicast (SSM) functionality and related configuration. However, the Firewall Threat Defense device allows SSM-related packets to pass through unless it is placed as a last-hop router.

SSM is classified as a data delivery mechanism for one-to-many applications such as IPTV. The SSM model uses a concept of "channels" denoted by an (S,G) pair, where S is a source address and G is an SSM destination address. Subscribing to a channel is achieved by using a group management protocol such as IGMPv3. SSM enables a receiving client, once it has learned about a particular multicast source, to receive multicast streams directly from the source rather than receiving it from a shared Rendezvous Point (RP). Access control mechanisms are introduced within SSM providing a security enhancement not available with current sparse or sparse-dense mode implementations.

PIM-SSM differs from PIM-SM in that it does not use an RP or shared trees. Instead, information on source addresses for a multicast group is provided by the receivers through the local receivership protocol (IGMPv3) and is used to directly build source-specific trees.

Multicast Bidirectional PIM

Multicast bidirectional PIM is useful for networks that have many sources and receivers talking to each other simultaneously and where each participant can become both the source and receiver of multicast traffic, such as in videoconferencing, Webex meetings, and group chat. When PIM bidirectional mode is used, the RP only creates the (*,G) entry for the shared tree. There is no (S,G) entry. This conserves resources on the RP because state tables for each (S,G) entry are not maintained.

In PIM sparse mode, traffic only flows down the shared tree. In PIM bidirectional mode, traffic flows up and down the shared tree.

PIM bidirectional mode also does not use the PIM register/register-stop mechanism to register sources to the RP. Each source can begin sending to the source at any time. When the multicast packets arrive at the RP, they are forwarded down the shared tree (if there are receivers) or dropped (when there are no receivers). However, there is no way for the RP to tell the source to stop sending multicast traffic.

Design-wise you must think about where to place the RP in your network because it should be somewhere in the middle between the sources and receivers in the network.

PIM bidirectional mode has no Reverse Path Forwarding (RPF) check. Instead it uses the concept of a Designated Forwarder (DF) to prevent loops. This DF is the only router on the segment that is allowed to send multicast traffic to the RP. If there is only one router per segment that forwards multicast traffic, there will be no loops. The DF is chosen using the following mechanism:

- The router with the lowest metric to the RP is the DF.
- If the metric is equal, then the router with the highest IP address becomes the DF.

PIM Bootstrap Router (BSR)

PIM Bootstrap Router (BSR) is a dynamic Rendezvous Point (RP) selection model that uses candidate routers for RP function and for relaying the RP information for a group. The RP function includes RP discovery and provides a default route to the RP. It does this by configuring a set of devices as candidate BSRs (C-BSR) which participate in a BSR election process to choose a BSR amongst themselves. Once the BSR is chosen, devices that are configured as candidate Rendezvous Points (C-RP) start sending their group mapping to the elected BSR. The BSR then distributes the group-to-RP mapping information to all the other devices down the multicast tree through BSR messages that travel from PIM router to PIM router on a per-hop basis.

This feature provides a means of dynamically learning RPs, which is very essential in large complex networks where an RP can periodically go down and come up.

PIM Bootstrap Router (BSR) Terminology

The following terms are frequently referenced in the PIM BSR configuration:

- Bootstrap Router (BSR) — A BSR advertises Rendezvous Point (RP) information to other routers with PIM on a hop-by-hop basis. Among multiple Candidate-BSRs, a single BSR is chosen after an election process. The primary purpose of this Bootstrap router is to collect all Candidate-RP (C-RP) announcements in to a database called the RP-set and to periodically send this out to all other routers in the network as BSR messages (every 60 seconds).
- Bootstrap Router (BSR) messages — BSR messages are multicast to the All-PIM-Routers group with a TTL of 1. All PIM neighbors that receive these messages retransmit them (again with a TTL of 1) out of all interfaces except the one in which the messages were received. BSR messages contain the RP-set and the IP address of the currently active BSR. This is how C-RPs know where to unicast their C-RP messages.
- Candidate Bootstrap Router (C-BSR) — A device that is configured as a candidate-BSR participates in the BSR election mechanism. A C-BSR with highest priority is elected as the BSR. The highest IP address of the C-BSR is used as a tiebreaker. The BSR election process is preemptive, for example if a new C-BSR with a higher priority comes up, it triggers a new election process.
- Candidate Rendezvous Point (C-RP) — An RP acts as a meeting place for sources and receivers of multicast data. A device that is configured as a C-RP periodically advertises the multicast group mapping information directly to the elected BSR through unicast. These messages contain the Group-range, C-RP address, and a hold time. The IP address of the current BSR is learned from the periodic BSR messages that are received by all routers in the network. In this way, the BSR learns about possible RPs that are currently up and reachable.

Note The Firewall Threat Defense device does not act as a C-RP, even though the C-RP is a mandatory requirement for BSR traffic. Only routers can act as a C-RP. So, for BSR testing functionality, you must add routers to the topology.

- BSR Election Mechanism — Each C-BSR originates Bootstrap messages (BSMs) that contain a BSR Priority field. Routers within the domain flood the BSMs throughout the domain. A C-BSR that hears about a higher-priority C-BSR than itself suppresses its sending of further BSMs for some period of time. The single remaining C-BSR becomes the elected BSR, and its BSMs inform all the other routers in the domain that it is the elected BSR.

Multicast Group Concept

Multicast is based on the concept of a group. An arbitrary group of receivers expresses an interest in receiving a particular data stream. This group does not have any physical or geographical boundaries—the hosts can be located anywhere on the Internet. Hosts that are interested in receiving data flowing to a particular group must join the group using IGMP. Hosts must be a member of the group to receive the data stream.

Multicast Addresses

Multicast addresses specify an arbitrary group of IP hosts that have joined the group and want to receive traffic sent to this group.

Clustering

Multicast routing supports clustering. In Spanned EtherChannel clustering, the control unit sends all multicast routing packets and data packets until fast-path forwarding is established. After fast-path forwarding is established, data units may forward multicast data packets. All data flows are full flows. Stub forwarding flows are also supported. Because only one unit receives multicast packets in Spanned EtherChannel clustering, redirection to the control unit is common.

Requirements and Prerequisites for Multicast Routing

Model Support

Threat Defense

Firewall Threat Defense Virtual

Supported Domains

Any

User Roles

Admin

Network Admin

Guidelines for Multicast Routing

Firewall Mode

Supported only in routed firewall mode. Transparent firewall mode is not supported.

IPv6

Does not support IPv6.

Multicast Group

The range of addresses between 224.0.0.0 and 224.0.0.255 is reserved for the use of routing protocols and other topology discovery or maintenance protocols, such as gateway discovery and group membership reporting. Hence, Internet multicast routing from address range 224.0.0/24 is not supported; IGMP group is not created when enabling multicast routing for the reserved addresses.

Clustering

In clustering, for IGMP and PIM, this feature is only supported on the primary unit.

Additional Guidelines

- You must configure an access control or prefilter rule on the inbound security zone to allow traffic to the multicast host, such as 224.1.2.3. However, you cannot specify a destination security zone for the rule, or it cannot be applied to multicast connections during initial connection validation.
- You cannot disable an interface with PIM configured on it. If you have configured PIM on the interface (see Configure PIM Protocol, on page 1313), disabling the multicast routing and PIM does not remove the PIM configuration. You must remove (delete) the PIM configuration to disable the interface.
- PIM/IGMP Multicast routing is not supported on interfaces in a traffic zone.
- Do not configure Firewall Threat Defense to simultaneously be a Rendezvous Point (RP) and a First Hop Router.
- HSRP standby IP address does not participate in PIM neighborship. Thus, if the RP router IP is routed through a HSRP standby IP address, the multicast routing does not work in Firewall Threat Defense. Hence for the multicast traffic to pass through successfully, ensure that the route for the RP address is not the HSRP standby IP address, instead, configure the route address to an interface IP address.
- For a device using virtual routing, you can configure multicast only for its global virtual router and not for its user-defined virtual router.

Configure IGMP Features

IP hosts use IGMP to report their group memberships to directly-connected multicast routers. IGMP is used to dynamically register individual hosts in a multicast group on a particular LAN. Hosts identify group memberships by sending IGMP messages to their local multicast router. Under IGMP, routers listen to IGMP messages and periodically send out queries to discover which groups are active or inactive on a particular subnet.

This section describes how to configure optional IGMP settings on a per-interface basis.

Procedure

Step 1 Enable Multicast Routing, on page 1309.

Step 2 Configure IGMP Protocol, on page 1309.

Step 3 Configure IGMP Access Groups, on page 1311.

Step 4 Configure IGMP Static Groups, on page 1311.

Step 5 Configure IGMP Join Groups, on page 1312.

Enable Multicast Routing

Enabling multicast routing on the Firewall Threat Defense device, enables IGMP and PIM on all interfaces by default. IGMP is used to learn whether members of a group are present on directly attached subnets. Hosts join multicast groups by sending IGMP report messages. PIM is used to maintain forwarding tables to forward multicast datagrams.

Note Only the UDP transport layer is supported for multicast routing.

The following list shows the maximum number of entries for specific multicast tables. Once these limits are reached, any new entries are discarded.

- MFIB—30,000
- IGMP Groups—30,000
- PIM Routes—72,000

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **IGMP**.

Step 3 Check the **Enable Multicast Routing** check box.

Checking this check box enables IP multicast routing on the device. Unchecking this check box disables IP multicast routing. By default, multicast is disabled. Enabling multicast routing enables multicast on all interfaces.

You can disable multicast on a per-interface basis. This is useful if you know that there are no multicast hosts on a specific interface and you want to prevent the Firewall Threat Defense device from sending host query messages on that interface.

Configure IGMP Protocol

You can configure IGMP parameters per interface, such as the forward interface, query messages, and time intervals.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **IGMP**.

Step 3 On **Protocol**, click **Add** or **Edit**.

Use the **Add IGMP parameters** dialog box to add new IGMP parameters to the Firewall Threat Defense device. Use the **Edit IGMP parameters** dialog box to change existing parameters.

Step 4 Configure the following options:

- **Interface**—From the drop-down list, choose the interface for which you want to configure IGMP protocol.
- **Enable IGMP**—Check the check box to enable IGMP.

  Note
  Disabling IGMP on specific interfaces is useful if you know that there are no multicast hosts on a specific interface and you want to prevent the device from sending host query messages on that interface.

- **Forward Interface**—From the drop-down list, choose the specific interface from which you want to forward IGMP messages.

  This configures the Secure Firewall Threat Defense device to act as an IGMP proxy agent and forward IGMP messages from hosts connected on one interface to an upstream multicast router on another interface.

- **Version**—Choose IGMP Version 1 or 2.

  By default, the Firewall Threat Defense device runs IGMP Version 2, which enables several additional features.

  Note
  All multicast routers on a subnet must support the same version of IGMP. The Firewall Threat Defense device does not automatically detect Version 1 routers and switch to Version 1. However, you can have a mix of IGMP Version 1 and 2 hosts on the subnet; the Firewall Threat Defense device running IGMP Version 2 works correctly when IGMP Version 1 hosts are present.

- **Query Interval**—The interval in seconds at which the designated router sends IGMP host-query messages. The range is 1 to 3600. The default is 125.

  Note
  If the Firewall Threat Defense device does not hear a query message on an interface for the specified timeout value, then the device becomes the designated router and starts sending the query messages.

- **Response Time**—The interval in seconds before the Firewall Threat Defense device deletes the group. The range is 1 to 25. The default is 10.

  If the Firewall Threat Defense device does not receive a response to a host query within this amount of time, it deletes the group.

- **Group Limit**—The maximum number of hosts that can join on an interface. The range is 1 to 500. The default is 500.

  You can limit the number of IGMP states resulting from IGMP membership reports on a per-interface basis. Membership reports exceeding the configured limits are not entered in the IGMP cache, and traffic for the excess membership reports is not forwarded

- **Query Timeout**—The period of time in seconds before which the Firewall Threat Defense device takes over as the requester for the interface after the previous requester has stopped. The range is 60 to 300. The default is 255.

Step 5 Click **OK** to save the IGMP protocol configuration.

Configure IGMP Access Groups

You can control access to multicast groups by using access control lists.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **Access Group**.

Step 3 On **Access Group**, click **Add** or **Edit**.

Use the **Add IGMP Access Group parameters** dialog box to add new IGMP access groups to the Access Group table. Use the **Edit IGMP Access Group parameters** dialog box to change existing parameters.

Step 4 Configure the following options:

a) From the **Interface** drop-down list, choose the interface with which the access group is associated. You cannot change the associated interface when you are editing an existing access group.

b) Click one of the following:

- **Standard Access List**— From the **Standard Access List** drop-down list, choose the standard ACL or click **Add** (+) to create a new standard ACL. See Configure Standard ACL Objects, on page 1374 for the procedure.
- **Extended Access List**—From the **Extended Access List** drop-down list, choose the extended ACL or click **Add** (+) to create a new extended ACL. See Configure Extended ACL Objects, on page 1370 for the procedure.

Step 5 Click **OK** to save the access group configuration.

Configure IGMP Static Groups

Sometimes a group member cannot report its membership in the group or there may be no members of a group on the network segment, but you still want multicast traffic for that group to be sent to that network segment. You can have multicast traffic for that group sent to the segment by configuring a statically joined IGMP group. With this method, the Firewall Threat Defense device does not accept the packets itself, but only forwards them. Therefore, this method allows fast switching. The outgoing interface appears in the IGMP cache, but this interface is not a member of the multicast group.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **IGMP**.

Step 3 On **Static Group**, click **Add** or **Edit**.

Use the **Add IGMP Static Group parameters** dialog box to statically assign a multicast group to an interface. Use the **Edit IGMP Static Group parameters** dialog box to change existing static group assignments.

Note
The IGMP Static Group enables PIM to send *Join* requests towards the sources or towards the Rendezvous Point (RP), provided, the firewall with this command is the PIM Designated Router (DR) on that interface where the command is applied.

Step 4 Configure the following options:

- From the **Interface** drop-down list, choose the interface to which you want to statically assign a multicast group. If you are editing an existing entry, you cannot change the value.
- From the **Multicast Groups** drop-down list, choose the multicast group to which you want to assign the interface, or click **Add** (+) to create a new multicast group. See Creating Network Objects for the procedure.

Step 5 Click **OK** to save the static group configuration.

Configure IGMP Join Groups

You can configure an interface to be a member of a multicast group. Configuring the Firewall Threat Defense device to join a multicast group causes upstream routers to maintain multicast routing table information for that group and keep the paths for that group active.

Note See Configure IGMP Static Groups, on page 1311 if you want to forward multicast packets for a specific group to an interface without the Firewall Threat Defense device accepting those packets as part of the group.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **IGMP**.

Step 3 On **Join Group**, click **Add** or **Edit**.

Use the **Add IGMP Join Group parameters** dialog box to configure the Firewall Threat Defense device to be a member of a multicast group. Use the **Edit IGMP Join Group parameters** dialog box to change existing parameters.

Note
The IGMP Join Group enables PIM to send *Join* requests towards the sources or towards the Rendezvous Point (RP), provided, the firewall with this command is the PIM Designated Router (DR) on that interface where the command is applied.

Step 4 Configure the following options:

- From the **Interface** drop-down list, choose the interface you want to be a member of a multicast group. If you are editing an existing entry, you cannot change the value.
- From the **Join Group** drop-down list, choose the multicast group to which you want to assign the interface, or click **Plus** to create a new multicast group. See Creating Network Objects for the procedure.

Configure PIM Features

Routers use PIM to maintain forwarding tables to use for forwarding multicast diagrams. When you enable multicast routing on the Secure Firewall Threat Defense device, PIM and IGMP are automatically enabled on all interfaces.

Note PIM is not supported with PAT. The PIM protocol does not use ports, and PAT only works with protocols that use ports.

This section describes how to configure optional PIM settings.

Procedure

Step 1 Configure PIM Protocol, on page 1313.
Step 2 Configure PIM Neighbor Filters, on page 1314.
Step 3 Configure PIM Bidirectional Neighbor Filters, on page 1315.
Step 4 Configure PIM Rendezvous Points, on page 1316.
Step 5 Configure PIM Route Trees, on page 1317.
Step 6 Configure PIM Request Filters, on page 1317.
Step 7 Configure Multicast Boundary Filters, on page 1320.

Configure PIM Protocol

You can enable or disable PIM on a specific interface.

You can also configure the Designated Router (DR) priority. The DR is responsible for sending PIM register, join, and prune messages to the RP. When there is more than one multicast router on a network segment, choosing the DR is based on the DR priority. If multiple devices have the same DR priority, then the device with the highest IP address becomes the DR. By default, the Firewall Threat Defense device has a DR priority of 1.

Router query messages are used to choose the PIM DR. The PIM DR is responsible for sending router query messages. By default, router query messages are sent every 30 seconds. Additionally, every 60 seconds, the Firewall Threat Defense device sends PIM join or prune messages.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Protocol**, click **Add** or **Edit**.

Use the **Add PIM parameters** dialog box to add new PIM parameters to the interface. Use the **Edit PIM parameters** dialog box to change existing parameters.

Step 4 Configure the following options:

- **Interface**—From the drop-down list, select the interface for which you want to configure PIM protocol.
- **Enable PIM**—Check the check box to enable PIM.
- **DR Priority**—The value for the DR for the selected interface. The router with the highest DR priority on the subnet becomes the designated router. Valid values range from 0 to 4294967294. The default DR priority is 1. Setting this value to 0 makes the Firewall Threat Defense device interface ineligible to become the designated router.
- **Hello Interval**—The interval in seconds at which the interface sends PIM hello messages. The range is 1 to 3600. The default is 30.
- **Join Prune Interval**—The interval in seconds at which the interface sends PIM join and prune advertisements. The range is 10 to 600. The default is 60.

Step 5 Click **OK** to save the PIM protocol configuration.

Configure PIM Neighbor Filters

You can define the routers that can become PIM neighbors. By filtering the routers that can become PIM neighbors, you can do the following:

- Prevent unauthorized routers from becoming PIM neighbors.
- Prevent attached stub routers from participating in PIM.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Neighbor Filter**, click **Add** or **Edit**.

Use the **Add PIM Neighbor Filter** dialog box to add new PIM neighbor filters to the interface. Use the **Edit PIM Neighbor Filter** dialog box to change existing parameters.

Step 4 Configure the following options:

- From the **Interface** drop-down list, choose the interface to which you want to add a PIM neighbor filter.
- **Standard Access List**— From the **Standard Access List** drop-down list, choose a standard ACL or click **Add** (+) to create a new standard ACL. See Configure Standard ACL Objects, on page 1374 for the procedure.

  Note
  Choosing **Allow** on the **Add Standard Access List Entry** dialog box lets the multicast group advertisements pass through the interface. Choosing **Block** prevents the specified multicast group advertisements from passing through the interface. When a multicast boundary is configured on an interface, all multicast traffic is prevented from passing through the interface unless permitted with a neighbor filter entry.

Step 5 Click **OK** to save the PIM neighbor filter configuration.

Configure PIM Bidirectional Neighbor Filters

A PIM bidirectional neighbor filter is an ACL that defines the neighbor devices that can participate in the Designated Forwarder (DF) election. If a PIM bidirectional neighbor filter is not configured for an interface, there are no restrictions. If a PIM bidirectional neighbor filter is configured, only those neighbors permitted by the ACL can participate in the DF election process.

Bidirectional PIM allows multicast routers to keep reduced state information. All of the multicast routers in a segment must be bidirectionally enabled to elect a DF.

When a PIM bidirectional neighbor filter is enabled, the routers that are permitted by the ACL are considered to be bidirectionally capable. Therefore, the following is true:

- If a permitted neighbor does not support bidirectional mode, then the DF election does not occur.
- If a denied neighbor supports bidirectional mode, then the DF election does not occur.
- If a denied neighbor does not support bidirectional mode, the DF election can occur.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Multicast Routing** > **PIM**.

Step 3 On **Bidirectional Neighbor Filter**, click **Add** or **Edit**.

Use the **Add PIM Bidirectional Neighbor Filter** dialog box to create ACL entries for the PIM bidirectional neighbor filter ACL. Use the **Edit PIM Bidirectional Neighbor Filter** dialog box to change existing parameters.

Step 4 Configure the following options:

- From the **Interface** drop-down list, select the interface to which you want to configure the PIM bidirectional neighbor filter ACL entry.

- **Standard Access List**— From the **Standard Access List** drop-down list, select a standard ACL or click **Add** (+) to create a new standard ACL. See Configure Standard ACL Objects, on page 1374 for the procedure.

  Note
  Choosing **Allow** on the **Add Standard Access List Entry** dialog box lets the specified devices participate in the DR election process. Choosing **Block** prevents the specified devices from participating in the DR election process.

Step 5 Click **OK** to save the PIM bidirectional neighbor filter configuration.

Configure PIM Rendezvous Points

You can configure the Firewall Threat Defense device to serve as a RP to more than one group. The group range specified in the ACL determines the PIM RP group mapping. If an ACL is not specified, then the RP for the group is applied to the entire multicast group range (224.0.0.0/4). See Multicast Bidirectional PIM, on page 1305 for more information about bidirectional PIM.

The following restrictions apply to RPs:

- You cannot use the same RP address twice.
- You cannot specify All Groups for more than one RP.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Rendezvous Points**, click **Add** or **Edit**.

Use the **Add Rendezvous Point** dialog box to create a new entry to the Rendezvous Point table. Use the **Edit Rendezvous Point** dialog box to change existing parameters.

Step 4 Configure the following options:

- From the **Rendezvous Point IP address** drop-down list, choose the IP address that you want to add as an RP or click **Add** (+) to create a new network object. See Creating Network Objects for the procedure.
- Check the **Use bi-directional forwarding** check box if the specified multicast groups are to operate in bidirectional mode. In bidirectional mode, if the Firewall Threat Defense device receives a multicast packet and has no directly connected members or PIM neighbors present, it sends a prune message back to the source.
- Click **Use this RP for all Multicast Groups** to use the specified RP for all multicast groups on the interface.
- Click the **Use this RP for all Multicast Groups as specified below** to designate the multicast groups to use with the specified RP and then from the **Standard Access List** drop-down list, choose a standard

ACL or click **Add** (+) to create a new standard ACL. See Configure Standard ACL Objects, on page 1374 for the procedure.

Step 5 Click **OK** to save the rendezvous point configuration.

Configure PIM Route Trees

By default, PIM leaf routers join the shortest-path tree immediately after the first packet arrives from a new source. This method reduces delay, but requires more memory than the shared tree. You can configure whether or not the Firewall Threat Defense device should join the shortest-path tree or use the shared tree, either for all multicast groups or only for specific multicast addresses.

The shortest-path tree is used for any group that is not specified in the Multicast Groups table. The Multicast Groups table displays the multicast groups to use with the shared tree. The table entries are processed from the top down. You can create an entry that includes a range of multicast groups, but excludes specific groups within that range by placing deny rules for the specific groups at the top of the table and the permit rule for the range of multicast groups below the deny statements.

Note This behavior is known as Shortest Path Switchover (SPT). We recommend that you always use the Shared Tree option.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Route Tree**, select the path for the route tree:

- Click **Shortest Path** to use the shortest-path tree for all multicast groups.
- Click **Shared Tree** to use the shared tree for all multicast groups.
- Click **Shared tree for below mentioned group** to designate the groups specified in the Multicast Groups table, and then from the **Standard Access List** drop-down list, select a standard ACL or click **Add** (+) to create a new standard ACL. See Configure Standard ACL Objects, on page 1374 for the procedure.

Step 4 Click **OK** to save the route tree configuration.

Configure PIM Request Filters

When the Firewall Threat Defense device is acting as an RP, you can restrict specific multicast sources from registering with it to prevent unauthorized sources from registering with the RP. You can define the multicast sources from which the Firewall Threat Defense device will accept PIM register messages.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Request Filter**, define the multicast sources that are allowed to register with the Firewall Threat Defense device when it acts as an RP:

- From the **Filter PIM register messages using:** drop-down list select **None**, **Access List**, or **Route Map**.
- If you choose **Access List** from the drop-down list, select an extended ACL or click **Add** (+) to create a new extended ACL. See Configure Extended ACL Objects, on page 1370 for the procedure.

  Note
  In the **Add Extended Access List Entry** dialog box, select **Allow** from the drop-down list o create a rule that allows the specified source of the specified multicast traffic to register with the Firewall Threat Defense device, or select **Block** to create a rule that prevents the specified source of the specified multicast traffic from registering with the device.

- If you choose **Route Map**, select a route map from the **Route Map** drop-down list, or click **Add** (+) to create a new route map. See Creating Network Objects for the procedure.

Step 4 Click **OK** to save the request filter configuration.

Configure the Secure Firewall Threat Defense Device as a Candidate Bootstrap Router

You can configure the Firewall Threat Defense device as a candidate BSR.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Bootstrap Router**, check the **Configure this FTD as a Candidate Bootstrap Router (C-BSR)** check box to perform the C-BSR setup.

a) From the **Interface** drop-down list, select the interface on the Firewall Threat Defense device from which the BSR address is derived to make it a candidate.

This interface must be enabled with PIM.

b) In the **Hash mask length** field, enter the length of a mask (32 bits maximum) that is to be ANDed with the group address before the hash function is called. All groups with the same seed hash (correspond) to the same RP. For example, if this value is 24, only the first 24 bits of the group addresses matter. This fact allows you to get one RP for multiple groups. The range is 0 to 32.

c) In the **Priority** field, enter the priority of the candidate BSR. The BSR with the larger priority is preferred. If the priority values are the same, the router with the larger IP address is the BSR. The range is 0 to 255. The default value is 0.

Step 4 (Optional) Click **Add** (+) to select an interface on which no PIM BSR messages will be sent or received in the **Configure this FTD as a Border Bootstrap Router (BSR)** section.

- From the **Interface** drop-down list, select the interface on which no PIM BSR messages will be sent or received.

  RP or BSR advertisements are filtered effectively isolating two domains of RP information exchange.

- Check the **Enable Border BSR** check box to enable BSR.

Step 5 Click **OK** to save the bootstrap router configuration.

Configure Multicast Routes

Configuring static multicast routes lets you separate multicast traffic from unicast traffic. For example, when a path between a source and destination does not support multicast routing, the solution is to configure two multicast devices with a GRE tunnel between them and to send the multicast packets over the tunnel.

When using PIM, the Firewall Threat Defense device expects to receive packets on the same interface where it sends unicast packets back to the source. In some cases, such as bypassing a route that does not support multicast routing, you may want unicast packets to take one path and multicast packets to take another.

Static multicast routes are not advertised or redistributed.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **Multicast Routes** > **Add or Edit**.

Use the **Add Multicast Route Configuration** dialog box to add a new multicast route to the Firewall Threat Defense device. Use the **Edit Multicast Route Configuration** dialog box to change an existing multicast route.

Step 3 From the **Source Network** drop-down box, choose an existing network or click **Add** (+) to add a new one. See Creating Network Objects for the procedure.

Step 4 To configure an interface to forward the route, click **Interface** and configure the following options:

- From the **Source Interface** drop-down list, choose the incoming interface for the multicast route.
- From the **Output Interface/Dense** drop-down list, choose the destination interface that the route is forwarded through.
- In the **Distance** field, enter the distance of the multicast route. The range is 0 to 255.

Step 5 To configure an RPF address to forward the route, click **Address** and configure the following options:

- In the **RPF Address** field, enter the IP address for the multicast route.
- In the **Distance** field, enter the distance of the multicast route The range is 0 to 255.

Step 6 Click **OK** to save the multicast routes configuration.

Configure Multicast Boundary Filters

Address scoping defines domain boundary filters so that domains with RPs that have the same IP address do not leak into each other. Scoping is performed on the subnet boundaries within large domains and on the boundaries between the domain and the Internet.

You can set up an administratively scoped boundary filter on an interface for multicast group addresses. IANA has designated the multicast address range from 239.0.0.0 to 239.255.255.255 as the administratively scoped addresses. This range of addresses can be reused in domains administered by different organizations. The addresses would be considered local, not globally unique.

A standard ACL defines the range of affected addresses. When a boundary filter is set up, no multicast data packets are allowed to flow across the boundary from either direction. The boundary filter allows the same multicast group address to be reused in different administrative domains.

You can configure, examine, and filter Auto-RP discovery and announcement messages at the administratively scoped boundary. Any Auto-RP group range announcements from the Auto-RP packets that are denied by the boundary ACL are removed. An Auto-RP group range announcement is permitted and passed by the boundary filter only if all addresses in the Auto-RP group range are permitted by the boundary ACL. If any address is not permitted, the entire group range is filtered and removed from the Auto-RP message before the Auto-RP message is forwarded.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **Multicast Boundary Filter**, and then click **Add** or **Edit**.

Use the **Add Multicast Boundary Filter** dialog box to add new multicast boundary filters to the device. Use the **Edit Multicast Boundary Filter** dialog box to change existing parameters.

You can configure a multicast boundary for administratively scoped multicast addresses. A multicast boundary restricts multicast data packet flows and enables reuse of the same multicast group address in different administrative domains. When a multicast boundary is defined on an interface, only the multicast traffic permitted by the filter ACL passes through the interface.

Step 3 From the **Interface** drop-down list, choose the interface for which you are configuring the multicast boundary filter ACL.

Step 4 From the **Standard Access List** drop-down list, choose the standard ACL you want to use, or click **Add** (+) to create a new standard ACL. See Configure Standard ACL Objects, on page 1374 for the procedure.

Step 5 Check the **Remove any Auto-RP group range announcement from the Auto-RP packets that are denied by the boundary** check box to filter Auto-RP messages from sources denied by the boundary ACL. If this check box is not checked, all Auto-RP messages are passed.

Step 6 Click **OK** to save the multicast boundary filter configuration.

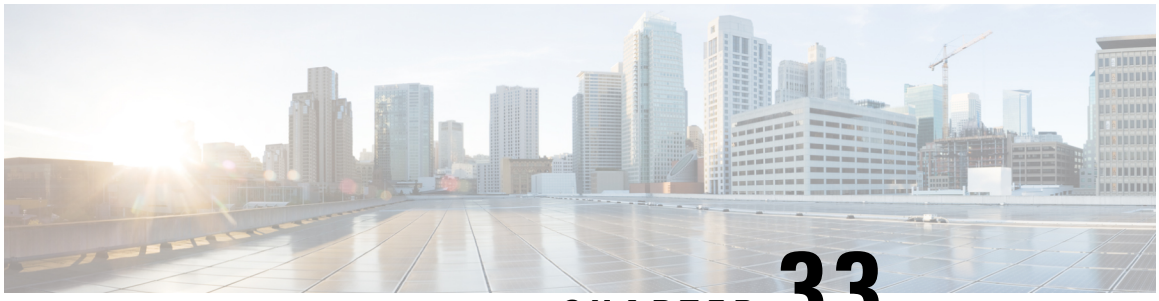

CHAPTER 33

Policy Based Routing

This chapter describes how to configure Firewall Threat Defense to support policy based routing (PBR) through Firewall Management Center's Policy based Routing page. The following sections describe policy based routing, guidelines for PBR, and configuration for PBR.

- About Policy Based Routing, on page 1323
- Guidelines and Limitations for Policy Based Routing, on page 1325
- Path Monitoring, on page 1326
- Configure Policy-Based Routing Policy, on page 1329
- Configuration Example for Policy Based Routing, on page 1339
- Configuration Example for PBR with Path Monitoring, on page 1344
- History for Policy Based Routing, on page 1346

About Policy Based Routing

In traditional routing, packets are routed based on the destination IP address. However, it is difficult to change the routing of specific traffic in a destination-based routing system. Policy Based Routing (PBR) gives you more control over routing by extending and complementing the existing mechanisms provided by routing protocols.

PBR allows you to set the IP precedence. It also allows you to specify a path for certain traffic, such as priority traffic over a high-cost link. With PBR, you can define routing that is based on criteria other than destination network such as source port, destination address, destination port, protocol, applications, or a combination of these objects.

You can use PBR to classify the network traffic based on applications, username, group membership, and security group association. This routing method is applicable in scenarios where, numerous devices access applications and data in a large network deployment. Traditionally, large deployments have topologies that backhaul all the network traffic to a hub as encrypted traffic in a route-based VPN. These topologies often result in issues such as packet latency, reduced bandwidth, and packet drop. Overcoming these issues involves high-cost complex deployments and management.

PBR policy enables you to securely breakout traffic for specified applications. You can configure PBR policy in the Secure Firewall Management Center user interface to allow the applications to be directly accessed.

Why Use Policy Based Routing

Consider a company that has two links between locations: one a high-bandwidth, low-delay expensive link, and the other a low-bandwidth, higher-delay, less-expensive link. While using traditional routing protocols,

the higher-bandwidth link gets most, if not all, of the traffic sent across it based on the metric savings obtained by the bandwidth, delay, or both (using EIGRP or OSPF) characteristics of the link. With PBR, you can route higher priority traffic over the high-bandwidth/low-delay link, while sending all other traffic over the low-bandwidth/high-delay link.

Following are a few scenarios where you can use Policy Based Routing:

Direct Internet Access

In this topology, application traffic from the branch office can be routed directly to the internet instead of through the VPN tunnel connecting to the headquarters. The branch Firewall Threat Defense is configured with an internet exit point and the PBR policy is applied on the ingress interface (*Inside 1*) to identify the traffic based on the applications, user identity (username and group membership), and Security Group Tag (security group association) defined in the ACL. Correspondingly, the traffic is forwarded through the egress interfaces directly to the internet or to the IPsec VPN tunnel.

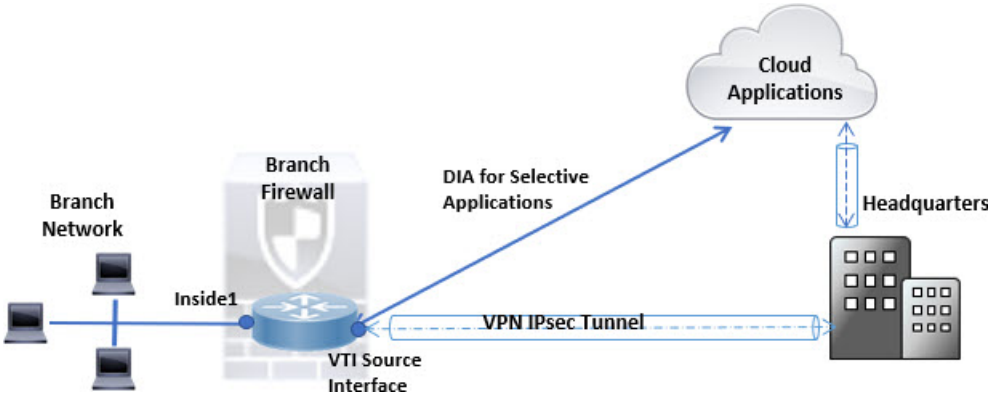

Equal-Access and Source-Sensitive Routing

In this topology, traffic from the HR and Mgmt networks can be configured to go through ISP1 and traffic from Eng network can be configured to go through ISP2. Thus, policy based routing enables the network administrators to provide equal-access and source-sensitive routing, as shown here.

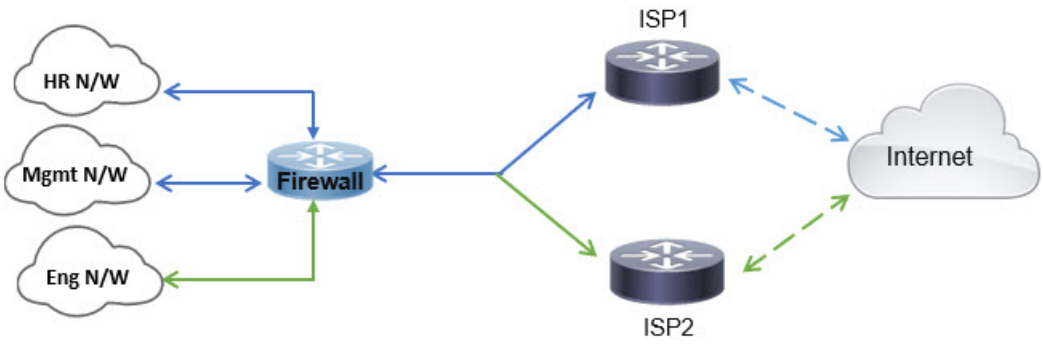

Load Sharing

In addition to the dynamic load-sharing capabilities offered by ECMP load balancing, network administrators can now implement policies to distribute traffic among multiple paths based on the traffic characteristics.

As an example, in the topology depicted in the Equal-Access Source Sensitive Routing scenario, an administrator can configure policy based routing to route the traffic from HR network through ISP1 and traffic from Eng network through ISP2 and thus share the load.

Guidelines and Limitations for Policy Based Routing

Firewall Mode Guidelines

PBR is supported only on routed firewall mode.

Device Guidelines

- PBR through Firewall Management Center's Policy Based Routing page is supported only from Version 7.1+ on both the Firewall Management Center and the device.
- When you upgrade Firewall Management Center or Firewall Threat Defense to version 7.1 and higher, the PBR configuration in the device is removed. You must configure PBR again using the Policy Based Routing page. If the managed device is lower than version 7.1, you must configure PBR again using FlexConfig with deploy option set to "every time."
- Configuring application, user identity, and Security Group Tag (SGT) based PBR policy on cluster devices is not supported.

Interface Guidelines

- Only routed interfaces and non management-only interfaces belonging to the Global virtual router can be configured as ingress or egress interface.
- PBR is not supported on user-defined virtual routers.
- Only interfaces that have a logical name can be defined in the policy.
- Static VTIs can be configured only as egress interfaces.
- Before proceeding with configuration, ensure that the ingress and egress traffic of each session flows through the same ISP-facing interface to avoid unexpected behavior caused by asymmetric routing, specifically when NAT and VPN are in use.

IPv6 Support

PBR supports IPv6.

Application-Based PBR and DNS Configuration

- Application-based PBR uses DNS snooping for application detection. Application detection succeeds only if the DNS requests pass through Firewall Threat Defense in a clear-text format; the DNS traffic is not encrypted.
- You must configure trusted DNS servers.

For more information on configuring DNS servers, see DNS, on page 949.

PBR Policies Not Applied for Output Route Look-up

Policy Based Routing is an ingress-only feature; that is, it is applied only to the first packet of a new incoming connection, at which time the egress interface for the forward leg of the connection is selected. Note that PBR

will not be triggered if the incoming packet belongs to an existing connection, or if NAT is applied and NAT chooses the egress interface.

PBR Policies Not Applied for Embryonic Traffic

Note An embryonic connection is where the necessary handshake between source and destination has not been made.

When a new internal interface is added and a new VPN policy is created using a unique address pool, PBR is applied to the outside interface matching the source of the new client pool. Thus, PBR sends traffic from the client to the next hop on the new interface. However, PBR is not involved in the return traffic from a host that has not yet established a connection with the new internal interface routes to the client. Thus, the return traffic from the host to the VPN client, specifically, the VPN client response is dropped as there is no valid route. You must configure a weighted static route with a higher metric on the internal interface.

HTTP-based Path Monitoring Guidelines

- HTTP-based path monitoring is supported on physical, port-channel, subinterfaces, and static tunnel interfaces. It is not supported on cluster devices.
- HTTP uses only IPv4 to ping the applications. IPv4 metrics are applied for routing and forwarding the IPv4 and IPv6 traffic.
- HTTP-based application monitoring is enabled by default Secure Firewall Management Center. However, when you upgrade from previous versions, this option is not enabled by default. You must manually enable it.

Additional Guidelines

- All existing configuration restrictions and limitations of route map will be carried forward.
- While defining the ACL for the policy match criteria, you can select multiple applications from a list of predefined applications to form an Access Control Entry (ACE). In Firewall Threat Defense, the predefined applications are stored as Network Service objects and the group of applications as Network Service Groups (NSG). You can create a maximum of 1024 such NSGs. The application or network service group is detected through first-packet classification. Currently, you cannot add to or modify the predefined applications list. However, you can create a custom application detector. See Create Custom Application Detector for PBR, on page 1332.
- Unicast Reverse Path Forwarding (uRPF) validates the source IP address of packets received on an interface against the routing table and not against the PBR route map. When uRPF is enabled, packets received on an interface through PBR are dropped as they are without the specific route entry. Hence, when using PBR, ensure to disable uRPF.

Path Monitoring

PBR uses either static cost or path monitoring (dynamic metrics) to route its traffic.

Path monitoring, when configured on interfaces, derive metrics such as round trip time (RTT), jitter, mean opinion score (MOS), and packet loss per interface. These metrics are used to determine the best path for routing PBR traffic.

ICMP-based Path Monitoring

The metrics on the interfaces are collected dynamically using ICMP probe messages to the interface's default gateway or a specified remote peer.

HTTP-based Path Monitoring

Path monitoring computes flexible metrics for multiple remote peers per interface. To monitor and determine best path for multiple applications through a policy on a branch firewall, HTTP is preferred over ICMP for the following reasons:

- HTTP-ping can derive the performance metrics of the path up to the application layer of the server, where the application is hosted.
- The need to change the firewall configuration whenever the application server IP address is changed is removed as the application domain is tracked instead of the IP address.

Note You can configure both ICMP and HTTP on the same interface. If the destination in the policy matches to any of the domain IP, the corresponding metrics are used. If the destination does not match to any of the configured domains, the metrics from the ICMP are used by PBR to select the outgoing interface.

Default Monitoring Timers

For metric collection and monitoring, the following timers are used:

- The interface monitor average interval is 30 seconds. This interval indicates the frequency to which the probes average.
- The interface monitor update interval is 30 seconds. This interval indicates the frequency at which the average of the collected values are calculated and made available for PBR to determine the best routing path.
- The interface monitor probe interval by ICMP is one second. This interval indicates the frequency at which an ICMP ping is sent.
- The application monitor probe interval by HTTP is 10 seconds. This interval indicates the frequency at which an HTTP ping is sent. Path monitoring uses the last 30 samples of HTTP ping for calculating the average metrics.

Note You cannot configure or modify the interval for any of these timers.

PBR and Path Monitoring

Typically, in PBR, traffic is forwarded through egress interfaces based on the priority value (interface cost) configured on them. From management center version 7.2, PBR uses IP-based path monitoring to collect the

performance metrics (RTT, jitter, packet-lost, and MOS) of the egress interfaces. PBR uses the metrics to determine the best path (egress interface) for forwarding the traffic. Path monitoring periodically notifies PBR about the monitored interface whose metric got changed. PBR retrieves the latest metric values for the monitored interfaces from the path monitoring database and updates the data path.

Path monitoring functions only with dynamic metrics, and only if the RTT, jitter, packet-lost, or MOS variables are set on the interfaces. Path monitoring does not function with static metrics—interface cost (cost set in interface).

You must enable path monitoring for the interface and configure the monitoring type. The PBR policy page allows you to specify the desired metric for path determination. See Configure Policy-Based Routing Policy, on page 1329.

PBR and HTTP-based Path Monitoring

From management center version 7.4, PBR can be configured to use HTTP-based path monitoring to collect the performance metrics of the application domains and not just one destination IP address. Path monitoring does not commence monitoring immediately after HTTP-based application monitoring is configured. It starts monitoring only when a DNS entry is snooped for a domain. With the information on the resolved IP for the domain, it sends and receives the HTTP request and response respectively. When DNS resolves multiple IP addresses for a single domain, the first resolved IP address will be used for probing and monitoring the application. It continues to monitor till the IP address changes or the HTTP-based path monitoring is disabled.

Based on the HTTP request and response durations, path monitoring computes the performance metrics for the application. The collected metrics are forwarded to PBR periodically, for making the routing and forwarding decision for the traffic arising from the configured ingress interface. If traffic arrives before the path monitoring could send its metrics to PBR, the traffic flow follows path chosen by the routing table. For the subsequent traffic flows that arrive after path monitoring metrics are available, PBR applies its routing decision based on the metrics and forwards the traffic.

Note Based on the Network Service Groups in the match ACL of the policy, you can apply PBR for multiple domains having multiple IP addresses.

In HTTP-based path monitoring of applications, the management center associates the applications/NSGs to the egress interface only when the PBR configuration meets the following criteria:

- The match ACL contains the monitored applications.
- The PBR policy is configured with any one of the following interface ordering values (metric type):
  - Minimal Jitter
  - Maximum Mean Opinion Score
  - Minimal Round-Trip Time
  - Minimal Packet Loss

Configure Path Monitoring Settings

The PBR policy relies on flexible metrics, such as round trip time (RTT), jitter, mean opinion score (MOS), and packet loss of the interfaces to identify the best routing path for its traffic. Path monitoring collects these

metrics on the specified interfaces. On the **Interfaces** page, you can configure interfaces with settings for path monitoring to send the ICMP probes or HTTP pings for metrics collection.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** (✎) for the interface you want to edit.

Step 3 Click the **Path Monitoring** tab.

Step 4 To configure ICMP based monitoring of the interface, click the **Enable IP based Monitoring** check box.

Step 5 From the **Monitoring Type** drop-down list, select the relevant option:

- **Auto**—Sends ICMP probes to the IPv4 default gateway of the interface. If the IPv4 gateway does not exist, path monitoring sends the probes to the IPv6 default gateway of the interface.
- **Peer IPv4**—Sends ICMP probes to the specified peer IPv4 address (next-hop IP) for monitoring. If you select this option, enter the IPv4 address in the **Peer IP To Monitor** field.
- **Peer IPv6**—Sends ICMP probes to the specified peer IPv6 address (next-hop IP) for monitoring. If you select this option, enter the IPv6 address in the **Peer IP To Monitor** field.
- **Auto IPv4**—Sends ICMP probes to the default IPv4 gateway of the interface.
- **Auto IPv6**—Send ICMP probes to the default IPv6 gateway of the interface.

Note

- The Auto options are not available for VTI interfaces. You must specify the peer address.
- Only one next-hop is monitored to a destination. That is, you cannot specify more than one peer address to monitor for an interface.

Step 6 By default, the **Enable HTTP based Application Monitoring** check box is selected. All applications selected for path monitoring in the match ACL of the PBR policy are listed, provided this interface is configured as the egress interface in the policy. To disable the HTTP-based monitoring of the interface, clear the check box.

Step 7 Click **Ok**, and to save the settings, click **Save**.

Configure Policy-Based Routing Policy

You can configure the PBR policy on the Policy Based Routing page by specifying the ingress interfaces, match criteria (Extended Access Control List), and egress interfaces.

Before you begin

To use the path monitoring metrics for configuring the traffic forwarding priority over egress interfaces, you must configure the path monitoring settings for the interfaces. See Configure Path Monitoring Settings, on page 1328.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Click **Routing**.

Step 3 Click **Policy Based Routing**.

The Policy Based Routing page displays the configured policy. The grid displays the list of ingress interfaces and a combination of the policy-based route access list, and egress interfaces.

Step 4 To configure the policy, click **Add**.

Step 5 In the **Add Policy Based Route** dialog box, select the **Ingress Interface** from the drop-down list.

Note
Only interfaces that have logical names and that belong to a global virtual router are listed in the drop-down. However, you cannot configure a VLAN interface with a logical name as the source (ingress) interface.

Step 6 To specify the match criteria and the forward action in the policy, click **Add**.

Step 7 In the **Add Forwarding Actions** dialog box, do the following:

a) From the **Match ACL** drop-down, choose the extended access control list object. You can predefine the ACL object (see Configure Extended ACL Objects, on page 1370) or click the **Add** (+) icon to create the object. In the **New Extended Access List Object** box, enter a name, click **Add** to open the **Add Extended Access List Entry** dialog box, where you can define the network, port, user identity, SGT, or application match criteria for the PBR policy. For information on creating custom domain to be synchronized with PBR, see Create Custom Application Detector for PBR, on page 1332.

Note
You can either have destination address or application/user identity/SGT defined in an ACE.

To selectively apply PBR on the incoming interface, you can define *Block* criteria in the ACE. When the traffic matches the block rule of the ACE, the traffic is forwarded to the egress interface based on the routing table.

b) From the **Send To** drop-down list:

- To select the configured interfaces, choose **Egress Interfaces**.
- To specify the IPv4/IPv6 next hop addresses, choose **IP Address**. Proceed to Step 7.e, on page 1331

c) If you have selected **Egress Interfaces**, from the **Interface Ordering** drop-down, choose the relevant option:

- By **Interface Priority**—The traffic is forwarded based on the priority of the interfaces. Traffic is routed to the interface with the least priority value first. When the interface is not available, the traffic is then forwarded to the interface with the next lowest priority value. For example, let us assume that *Gig0/1*, *Gig0/2*, and *Gig0/3* are configured with priority values *0*,*1*, and *2* respectively. The traffic is forwarded to *Gig0/1*. If *Gig0/1* becomes unavailable, the traffic is then forwarded to *Gig0/2*.

  Note
  To configure the priority for the interfaces, click **Configure Interface Priority** on the Policy Based Routing page. In the dialog box, provide the priority number against the interfaces, and then click **Save**. You can also configure the priority for an interface in the Interface Settings.

When the priority value is the same for all the interfaces, the traffic is balanced among the interfaces.

- By **Order**—The traffic is forwarded based on the sequence of the interfaces specified here. For example, let us assume that *Gig0/1*, *Gig0/2*, and *Gig0/3* are selected in the following order, *Gig0/2*, *Gig0/3*, *Gig0/1*. The traffic is forwarded to *Gig0/2* first, then to *Gig0/3*, irrespective of their priority values.
- By **Minimal Jitter**—The traffic is forwarded to the interface that has the lowest jitter value. You need to enable Path Monitoring on the interfaces for PBR to obtain the jitter values.
- By **Maximum Mean Opinion Score**—The traffic is forwarded to the interface that has the maximum mean opinion score (MOS). You need to enable Path Monitoring on the interfaces for PBR to obtain the MOS values.
- By **Minimal Round Trip Time**—The traffic is forwarded to the interface that has the minimal round trip time (RTT). You need to enable Path Monitoring on the interfaces for PBR to obtain the RTT values.
- By **Minimal Packet Loss**—The traffic is forwarded to the interface that has the minimal packet loss. You need to enable Path Monitoring on the interfaces for PBR to obtain the packet loss values.

d) In the **Available Interfaces** box, all the interfaces with their priority values are listed. From the list of interfaces, click the **Add** (+) button to add to the selected egress interfaces. Proceed to Step 7.k, on page 1332

Note
The route to the selected interface must be present in the routing table.

e) If you have selected **IP Address**, enter the IP addresses separated by commas in the **IPv4 Addresses** or **IPv6 Addresses** fields. The traffic is forwarded as per the sequence of the specified IP addresses.

Note
When multiple next-hop IP addresses are provided, the traffic is forwarded as per the sequence of the specified IP addresses until a valid routable next-hop IP address is found. The configured next-hops should be directly connected.

f) From the **Don't Fragment** drop-down list, select Yes, No, or None. If the DF (Don't Fragment) flag is set to *Yes*, the intermediate routers never perform fragmentation of a packet.

g) To specify the current interface as the default for forwarding, check the **Default Interface** check box.

h) The **IPv4 Settings** and **IPv6 Settings** tab enables you to specify the recursive and default settings:

Note
For a route-map, you can only specify either IPv4 or IPv6 next-hop settings.

- **Recursive**—The route map configuration is applied only when the specified next-hop address and the default next-hop address are found on a directly connected subnet. However, you could use the recursive option, where the next-hop address need not be directly connected. Here, a recursive lookup is performed on the next-hop address, and matching traffic is forwarded to the next-hop used by that route entry according to the current routing path of the router.
- **Default**—If the normal route lookup fails to match traffic, the traffic is forwarded to this specified next-hop IP address.

i) Check the **Peer Address** check box to use the next-hop address as the peer address.

Note
You cannot configure a route map with both default next-hop address and peer address.

j) For IPv4 settings, you can check whether the next IPv4 hops of a route map are available under **Verify Availability**—click the **Add** (+) button and add the next-hop IP address entries:

- **IP Address**—Enter the next hop IP address.
- **Sequence**—Entries are assessed in order using the sequence number. Ensure that no duplicate sequence numbers are entered. The valid range is 1 to 65535.
- **Track**—Enter a valid ID. The valid range is 1 to 255.

k) Click **Save**.

Step 8 To save the policy, click **Save** and **Deploy**.

The Firewall Threat Defense uses ACLs to match traffic and perform routing actions on the traffic. Typically, you configure a route map that specifies an ACL against which traffic is matched, and then you specify one or more actions for that traffic. With the use of path monitoring, PBR can now select the best egress interface for routing the traffic. Finally, you associate the route map with an interface on which you want to apply PBR on all incoming traffic.

Add Path Monitoring Dashboard

To view the path monitoring metrics, you must add the path monitoring dashboard to the Health Monitoring page of the device.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Step 2 Select the device, and click **Add New Dashboard**.

Step 3 Enter a name for the custom dashboard.

Step 4 In the **Metrics** area, click the **Add from Predefined Correlations** button.

Step 5 From the list, click **Interface - Path Metrics**.

By default, all the four metrics are selected for displaying as portlets in the dashboard with an additional metric field. You can exclude any of them by clicking **Delete** (🗑).

Step 6 Click **Add Dashboard**.

Create Custom Application Detector for PBR

This section outlines the process involved in creating custom application detectors and configuring application-based PBR policy.

1. Create a custom detector for a user-defined application, as described in Create a User-Defined Application, on page 2633.

Note While creating a user-defined application, ensure that you select NSG from the **Tag** drop-down list.

2. To manually define the detection pattern, select **Basic** detector type and proceed to define them, as described in Specifying Detection Patterns in Basic Detectors, on page 2634. Alternatively, if you want to create the detection pattern using a .lua file, select **Advanced** detector type and follow the instructions provided in Specifying Detection Criteria in Advanced Detectors, on page 2635.
3. Activate the custom detector.
4. Configure ACL policy (Extended) for PBR with the custom application ACE. For procedure to create the ACL, see Configure Extended ACL Objects, on page 1370.
5. In the PBR policy, select the ACL to be matched for defining forwarding actions. For detailed procedure to create PBR policy, see Configure Policy-Based Routing Policy, on page 1329.

Configuration Example for PBR with Custom Application Detector

This example details the configuration of PBR policy with custom application detectors for the following domain:

- amazon123.com
- flipkart.com
- hamleysonline.com

Before you begin

- This example assumes that you are aware of the basic steps to configure PBR policy and custom application detectors.
- You should have configured the ingress and egress interfaces with logical names. In this example, the ingress interface is named *Inside1*, and the egress interface is named *eBuy*.

Procedure

Step 1 Create a custom detector, *PBRePurchase*:

a) Choose **Policies** > **Application Detectors**, and click **Create Custom Detector**.
b) In the **Create A Custom Application Detector** field, enter a name (in this example, *PBRePurchase*) and a description.
c) To create a custom application, click (+).
d) In the the **Application Editor** dialog box, enter the relevant values in the fields. See Create a User-Defined Application, on page 2633 for a detailed description of each field.

Note
For the custom application to be detectable for PBR, choose **NSG** from the **Tag** drop-down list:

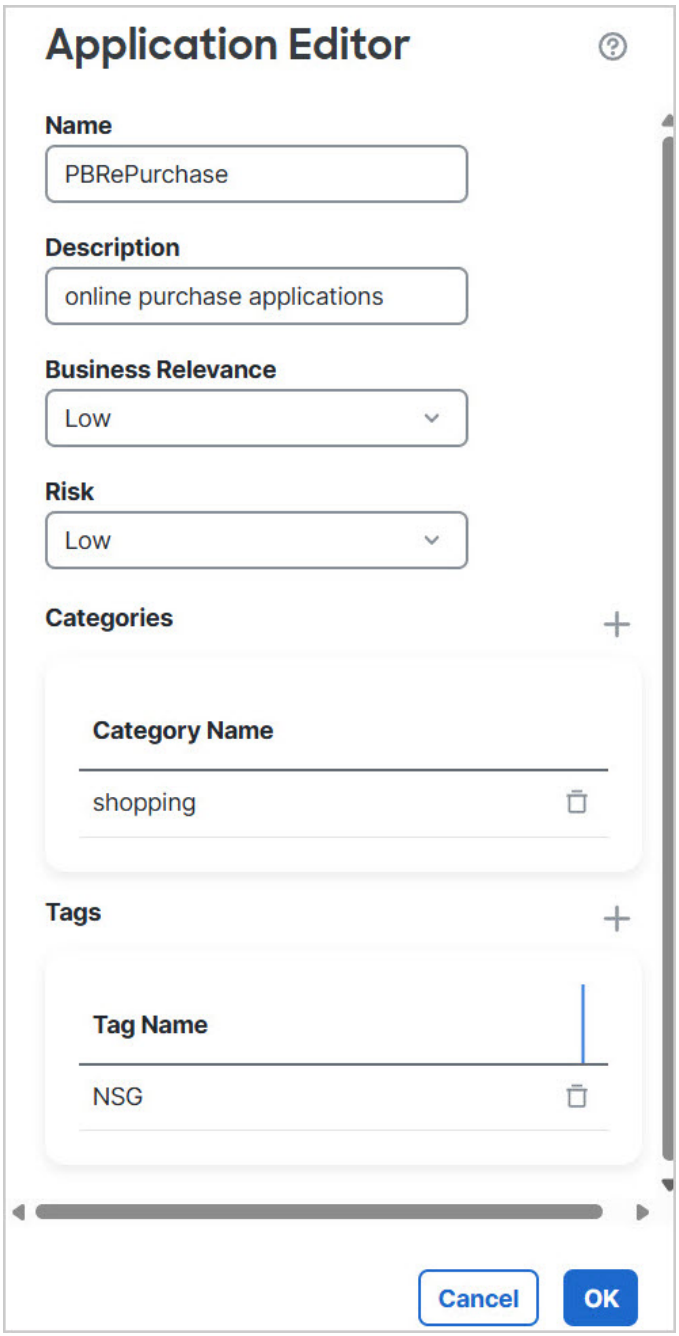

e) Click **Ok**.
f) From the **Application** drop-down list, choose *PBRePurchase*.

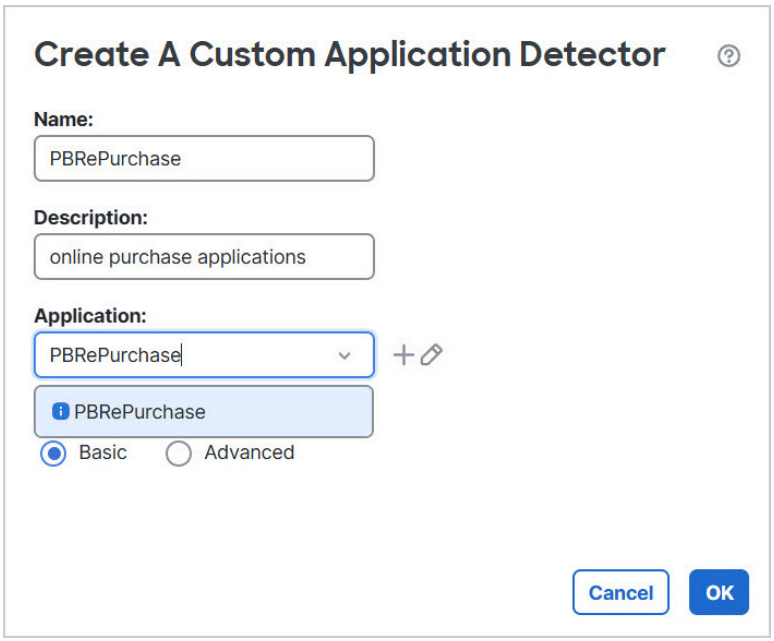

a) Click the **Basic** radio button, and then click **Ok**.

Note
The **Advanced** detector type is not supported in the PBR policy.

b) In the **Application Detector** page, click the **Add** button in the **Detection Patterns** area to add patterns.
c) From the **Application** drop-down list, choose **SSL** as the protocol type, and select the appropriate pattern type. Enter a pattern string that matches the selected type (in this example, enter *amazon123.com*), and then click **Ok**.

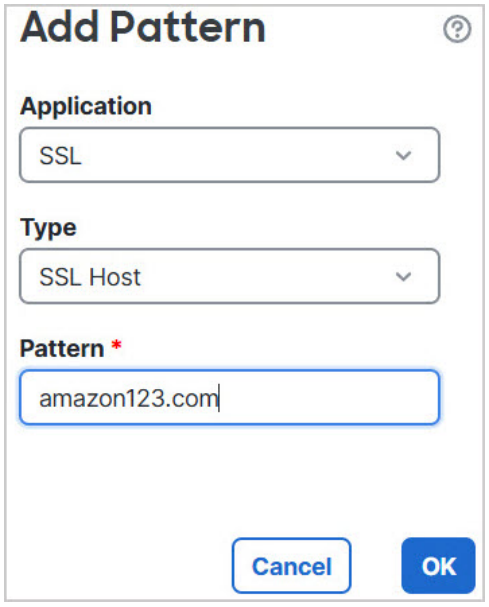

d) Repeat the procedure and create two more patterns for the custom domains *flipkart.com* and *hamleysonline.com*.

e) Click **Save**.

Step 2 In the application detector dashboard, use filters to search for the custom application detector that you have just created.

Step 3 To enable the custom detector, click the **State** toggle button ( )adjacent to the corresponding application detector.

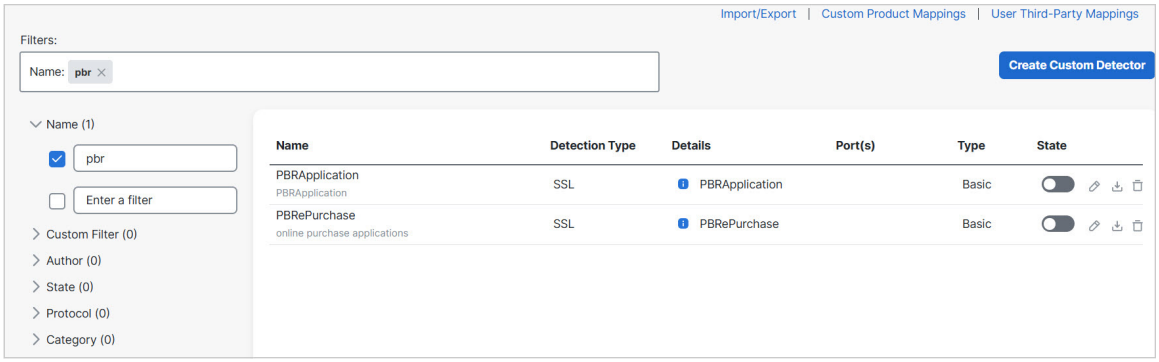

Step 4 In the dialog box that is displayed, click **Yes**.

A success message will appear when the application detector is activated:

Step 5 To create an ACL using the custom application detector for synchronizing with the PBR policy:

a) Choose **Objects** > **Object Management** > **Access List** > **Extended**.

b) Click **Add Extended Access List**.

c) Enter a name (for example, *PBR_shopping*) to the list and click **Add** to create an ACE for the list.

d) In the **Add Extended Access List Entry** dialog box, click the **Application** tab, select the application name *PBRePurchase*, and then click **Add to Rule**.

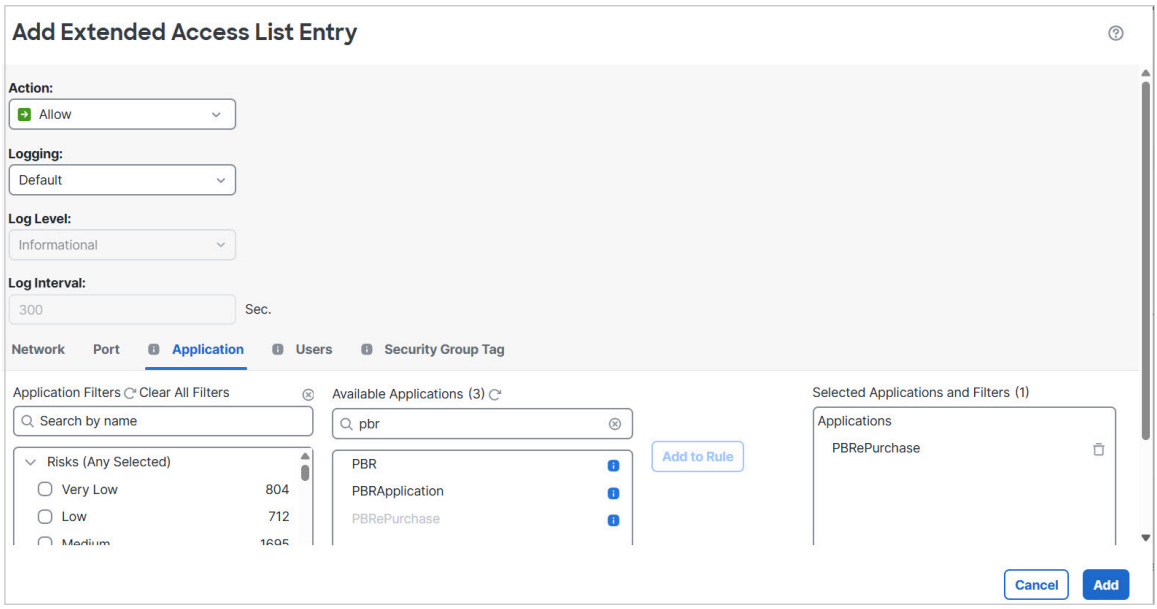

e) Click **Add**.

f) Click **Save**.

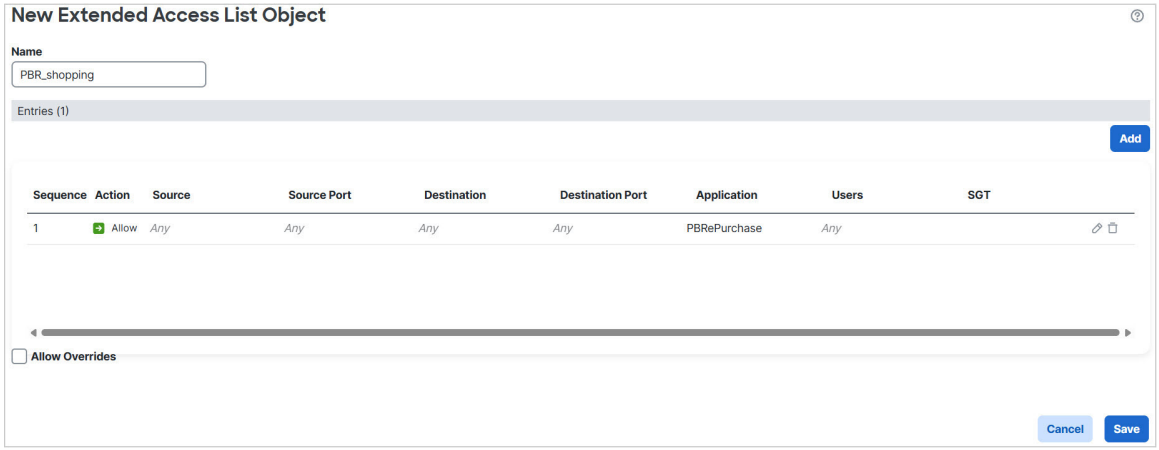

Step 6 Choose **Routing** > **Policy Based Routing**.

Step 7 On the **Policy Based Routing** page that is displayed, click **Add**.

Step 8 In the **Add Policy Based Route** dialog box, choose *Inside 1* from the **Ingress Interface** drop-down list.

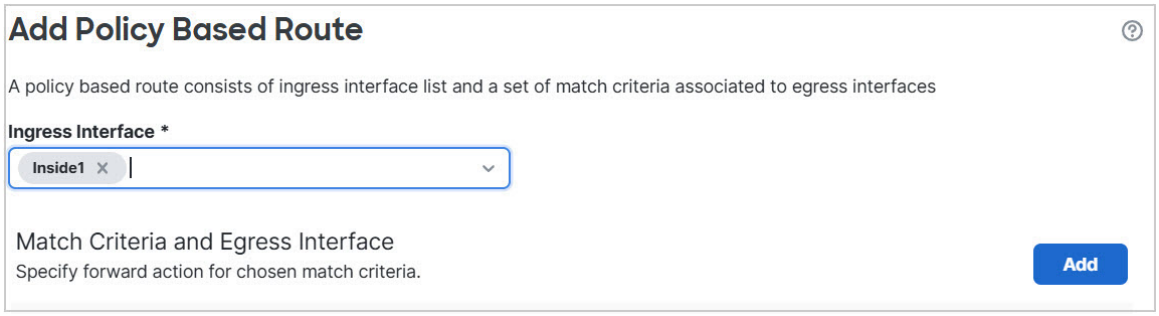

Step 9 From the **Match ACL** drop-down list, choose the ACL that you have created, which, in this example is *PBR_shopping* .

Note
In a Firewall Threat Defense, the application group in an ACL is configured as a network service group. That is why we tag the custom application with **NSG**.

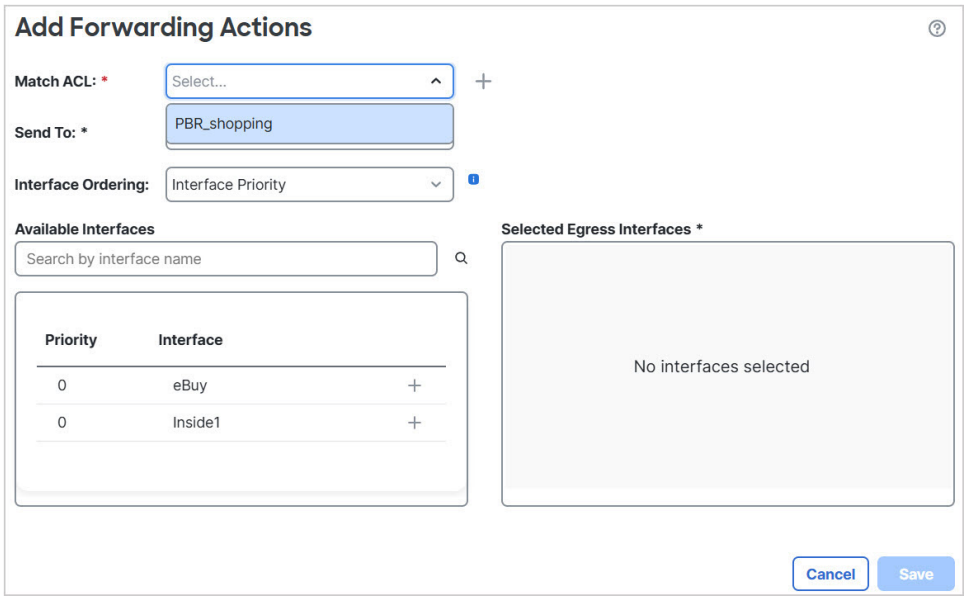

Step 10 Specify the egress interfaces:

a) From the **Send To** drop-down list, choose **Egress Interfaces**.
b) From the **Interface Ordering** drop-down list, choose the appropriate order.
c) Under **Available Interfaces**, click (+) adjacent to the corresponding interface, namely, *eBuy*.

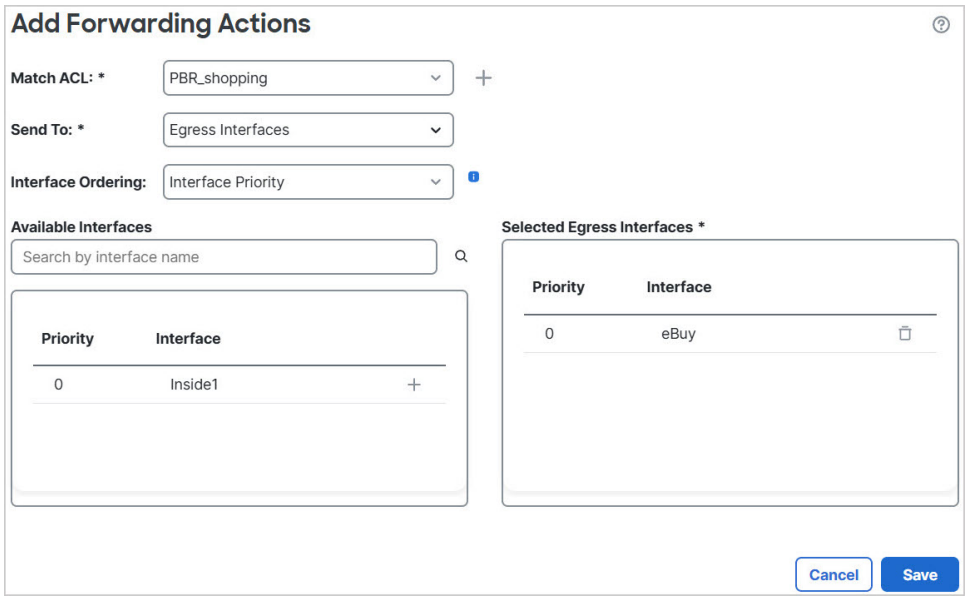

d) Click **Save**.

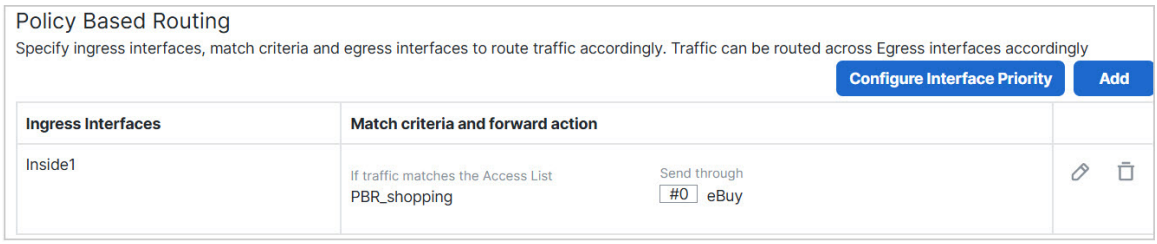

Step 11 **Save** and **Deploy**.

Configuration Example for Policy Based Routing

Consider a typical corporate network scenario where all the branch network traffic passes through a route-based VPN of the corporate network and diverges to the extranet, when required. Accessing the web-based applications that address day-to-day operations through the corporate network results in huge network expansion and maintenance costs. This example illustrates the PBR configuration procedure for direct internet access.

The following figure depicts the topology of a corporate network. The branch network is connected to the corporate network through a route-based VPN. Traditionally, the corporate Firewall Threat Defense is configured to handle both the internal and external traffic of the branch office. With the PBR policy, the branch Firewall Threat Defense is configured with a policy that routes specific traffic to the WAN network instead of the virtual tunnels. The rest of the traffic flows through the route-based VPN, as usual.

This example also illustrates the configuring of the WAN and the VTI interfaces with ECMP zones to achieve load balancing.

Figure 421: Configuring Policy Based Routing on Branch Firewall Threat Defense in Firewall Management Center

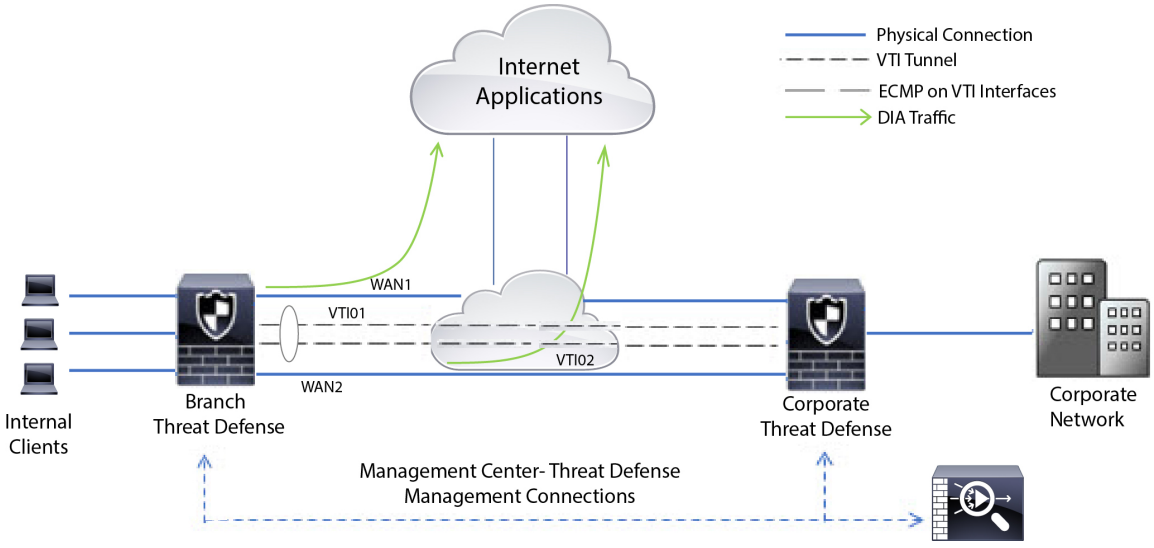

Before you begin

This example assumes that you have already configured WAN and VTI interfaces for the branch Firewall Threat Defense in Firewall Management Center.

Procedure

Step 1 Configure policy based routing for the branch Firewall Threat Defense, select the ingress interfaces:

a) Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

b) Choose **Routing** > **Policy Based Routing**, and on the **Policy Based Routing** page, click **Add**.

c) In the **Add Policy Based Route** dialog box, select the interfaces (say, *Inside 1*, and *Inside 2*) from the **Ingress Interface** drop-down list.

Step 2 Specify the match criteria:

a) Click **Add**.

b) To define the match criteria, click the **Add** (+) button.

c) In **New Extended Access List Object**, enter the name for the ACL (say, *DIA-FTD-Branch*), and click **Add**.

d) In the **Add Extended Access List Entry** dialog box, choose the required web-based applications from the **Application** tab:

Figure 422: Applications Tab

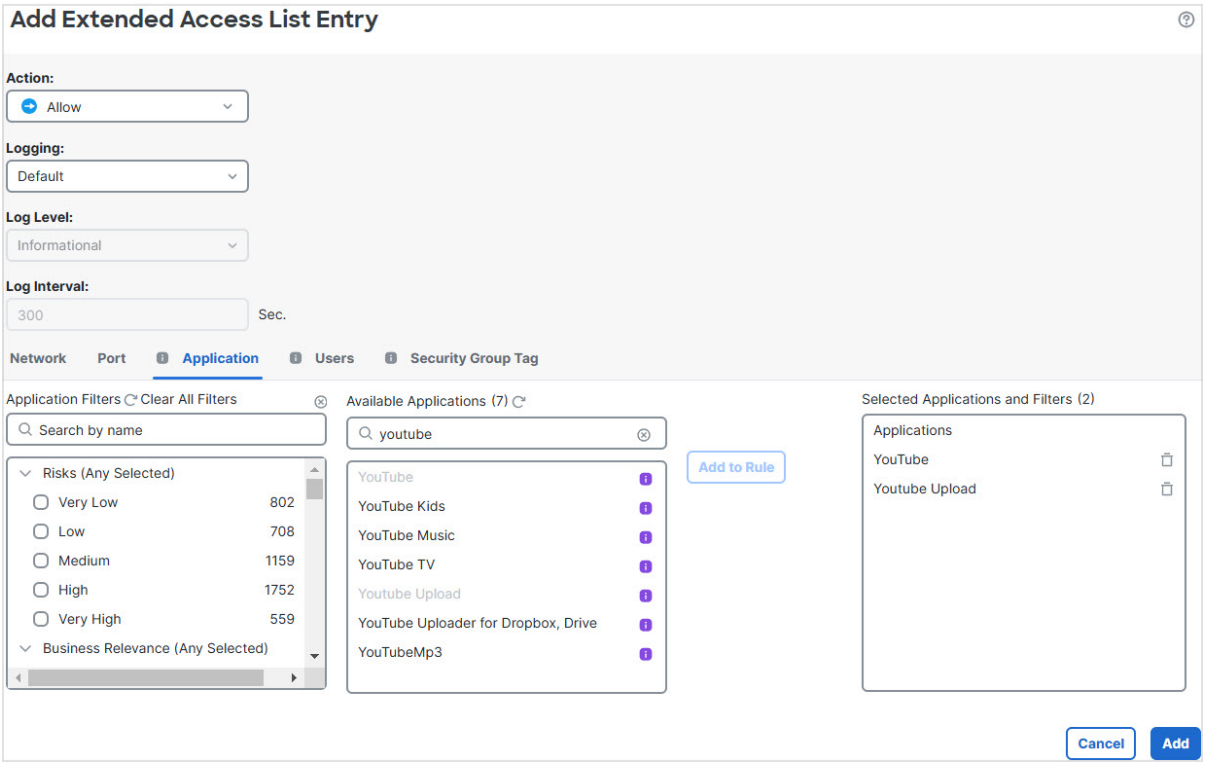

On the Firewall Threat Defense, the application group in an ACL is configured as a network service group and each of the applications as a network service object.

Figure 423: Extended ACL

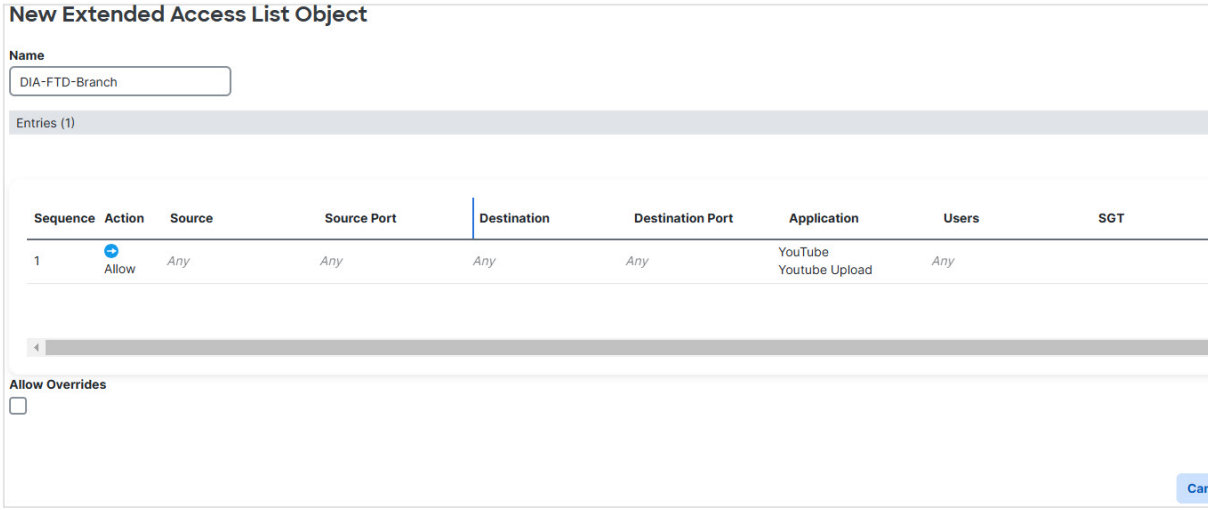

e) Click **Save**.

f) Select *DIA-FTD-Branch* from the **Match ACL** drop-down list.

Step 3 Specify the egress interfaces:

a) From the **Send To** and **Interface Ordering** drop-down lists, choose Egress Interfaces, and Interface Priority respectively.
b) Under **Available Interfaces**, click the ⊕ button against the respective interface names to add *WAN1* and *WAN2*:

Figure 424: Configuring Policy Based Routing

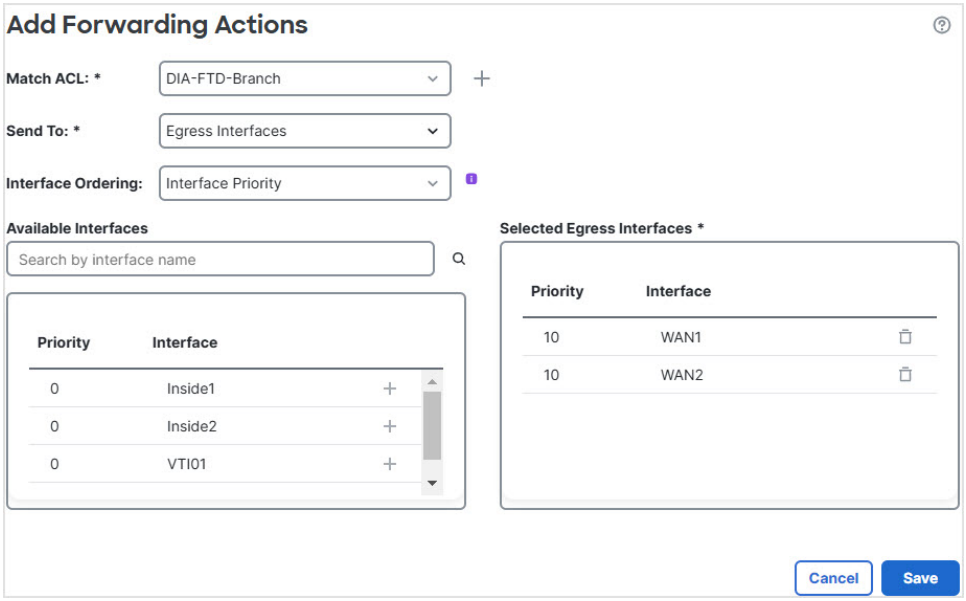

c) Click **Save**.

Step 4 Interface priority configuration:

You can set the priority value for the interfaces either in the **Edit Physical Interface** page, or in the **Policy Based Routing** page (**Configure Interface Priority**). In this example, the Edit Physical Interface method is described.

a) Choose **Devices** > **Device Management**, and edit the branch Firewall Threat Defense.
b) Set the priority for the interfaces. Click **Edit** against the interface and enter the priority value:

Figure 425: Setting Interface Priority

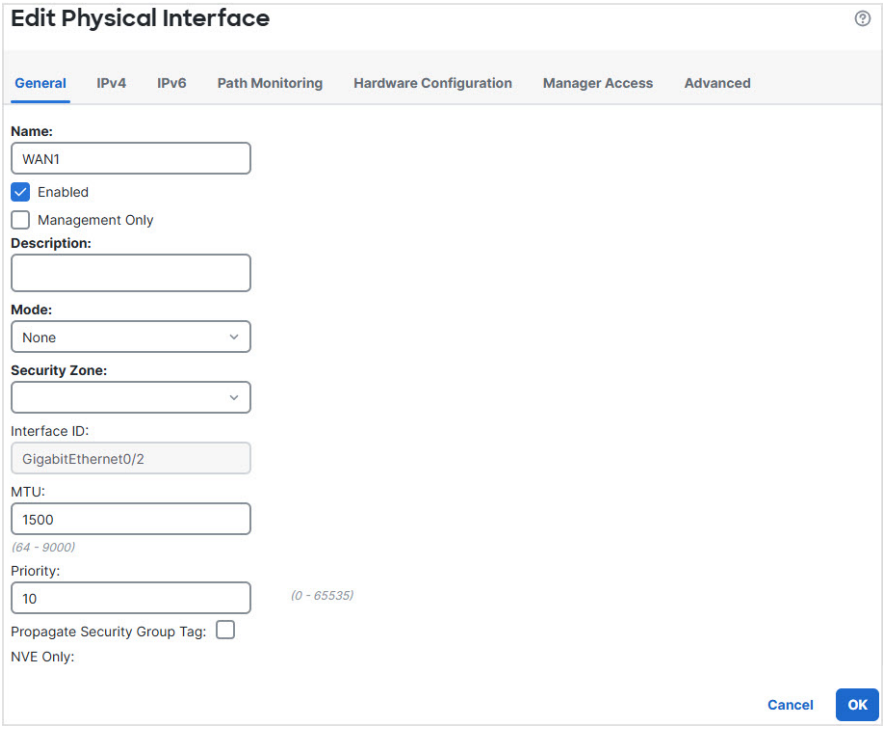

c) Click **Ok** and **Save**.

Step 5 Create ECMP zones for load balancing:

a) In the **Routing** page, click **ECMP**.

b) To associate interfaces to the ECMP zone, click **Add**.

c) Select *WAN1* and *WAN 2* and create an ECMP zone—*ECMP-WAN*. Similarly, add *VTI01* and *VTI02* and create an ECMP zone—*ECMP-VTI*:

Figure 426: Associating Interfaces with ECMP Zone

Step 6 Configure static routes for the zone interfaces for load balancing:

a) In the **Routing** page, click **Static Route**.

b) Click **Add** and specify the static routes for *WAN1*, *WAN2*, *VTI01*, and *VTI02*. Ensure that you specify the same metric value for the interfaces belonging to the same ECMP zones (Step 5):

Figure 427: Configuring Static Routes for ECMP Zone Interfaces

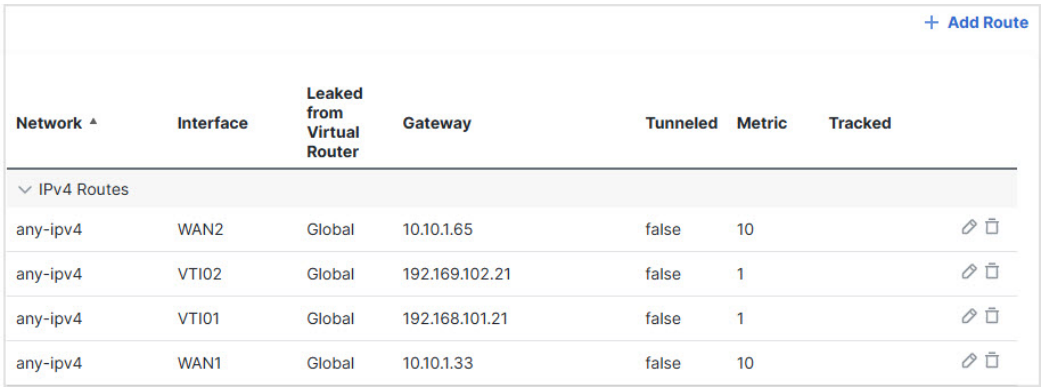

Note
Ensure that the zone interfaces have the same destination address and metric, but different gateway addresses.

Step 7 Configure trusted DNS on the WAN objects of the branch Firewall Threat Defense to ensure secured flow of traffic to the internet:

a) Choose **Devices** > **Platform Settings**, and create a DNS policy on the branch Firewall Threat Defense.
b) To specify the trusted DNS, **Edit** the policy and click **DNS**.
c) To specify the DNS servers for the DNS resolution to be used by WAN objects, in the **DNS Settings** tab, provide the DNS server group details and select WAN from the interface objects.
d) Use the **Trusted DNS Servers** tab to provide specific DNS servers that you trust for the DNS resolution.

Step 8 **Save** and **Deploy**.

Any *YouTube* related access requests from the branch inside network *INSIDE1* or *INSIDE2* are routed to *WAN1* or *WAN2* as they would match the *DIA-FTD-Branch* ACL. Any other request, say *google.com*, are routed through *VTI01* or *VTI02* as configured in the Site to Site VPN Settings.

With the ECMP configured, the network traffic is seamlessly balanced.

Configuration Example for PBR with Path Monitoring

This example details the configuration of PBR with path monitoring for the following applications with flexible metrics:

- Audio or video sensitive applications (example, WebEx Meetings) with Jitter.
- Cloud-based application (example, Office365) with RTT.
- Network-based access control (with a specific source and destination) with Packet Loss.

Before you begin

1. This example assumes that you are aware of the basic configuration steps for PBR.

2. You have configured ingress and egress interfaces with logical names. In this example, the ingress interface is named *Inside1*, and egress interfaces are named *ISP01*, *ISP02*, and *ISP03*.

Procedure

Step 1 Path monitoring configuration on interfaces *ISP01*, *ISP02*, and *ISP03*:

For the metrics collection on the egress interfaces, you must enable and configure path monitoring on them.

a) Choose **Devices** > **Device Management**, and edit Firewall Threat Defense.
b) Under the **Interfaces** tab, edit the interface (in our example, *ISP01*)
c) Click the **Path Monitoring** tab, select the **Enable Path Monitoring** check box, and then specify the monitoring type (see Configure Path Monitoring Settings, on page 1328).
d) Click **Ok** and **Save**.
e) Repeat the same steps and configure the path monitoring settings for *ISP02* and *ISP03*.

Step 2 Configure policy-based routing for a branch in an organization Firewall Threat Defense, select the ingress interfaces:

a) Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
b) Choose **Routing** > **Policy Based Routing**, and on the **Policy Based Routing** page, click **Add**.
c) In the **Add Policy Based Route** dialog box, select *Inside 1* from the **Ingress Interface** drop-down list.

Step 3 Specify the match criteria:

a) Click **Add**.
b) To define the match criteria, click the **Add** (+) button.
c) In **New Extended Access List Object**, enter the name for the ACL (example, *PBR-WebEx*), and click **Add**.
d) In the **Add Extended Access List Entry** dialog box, choose the required web-based applications (example, WebEx Meetings) from the **Application** tab.

Remember
On Firewall Threat Defense, the application group in an ACL is configured as a network service group and each of the applications as a network service object.

e) Click **Save**.
f) Select *PBR-WebEx* from the **Match ACL** drop-down list.

Step 4 Specify the egress interfaces:

a) From the **Send To** drop-down list, choose Egress Interfaces.
b) From the **Interface Ordering** drop-down list, choose By Minimal Jitter.
c) Under **Available Interfaces**, click the **Right Arrow** (>)button against the respective interface names to add *ISP01*, *ISP02*, and *ISP03*.
d) Click **Save**.

Step 5 Repeat Step 2 and Step 3 to create PBRs for the same interface (*Inside1*) to route Office365 and network-based access control traffic:

a) Create a match criteria object, example *PBR-Office365*, and select the Office365 application from the **Application** tab.
b) From the **Interface Ordering** drop-down list, choose By Minimal Round Trip Time.

c) Specify the egress interfaces *ISP01*, *ISP02*, and *ISP03*, and click **Save**.
d) Now, create a match criteria object, example *PBR-networks*, and specify the source and destination interface in the **Network** tab.
e) From the **Interface Ordering** drop-down list, choose By Minimal Packet Loss.
f) Specify the egress interfaces *ISP01*, *ISP02*, and *ISP03*, and click **Save**.

Step 6 **Save** and **Deploy**.

Step 7 To view path monitoring metrics, choose **Devices** > **Device Management**, and from **More** (⋮) click **Health Monitor**. To view the metric details for the interfaces of the device, you must add the path metrics dashboard. For details, see Add Path Monitoring Dashboard , on page 1332.

The WebEx, Office365, and networks-based ACL traffic are forwarded through the best route derived from the metrics value collected on *ISP01*, *ISP02*, and *ISP03*.

History for Policy Based Routing

Table 72:

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Custom Application Detector support in PBR	7.7	7.7	You can now create PBR polices with custom application domains. User-defined domains are created as custom detector patterns and can be used in Extended ACL for the application based PBR policy. New/modified screens: No new or modified screens were added.
Identity and SGT based PBR Policy	7.4.0	7.4.0	You can now classify the network traffic based on users and user groups, and SGTs in PBR policies. You can select the identity and SGT objects while defining the extended ACLs for the PBR policies. New/modified screens: New tabs added in the extended access list object for configuring the policy based routing policy: **Objects** > **Object Management** > **Access Control Lists** > **Add Extended** page, **Users** and **Security Group Tag**.
HTTP based Path Monitoring	7.4.0	7.2.0	PBR can now use the performance metrics (RTT, jitter, packet-lost, and MOS) collected by path monitoring through HTTP client on the application domain rather than the metrics on a specific destination IP. HTTP-based application monitoring option is enabled by default for the interface. You can configure a PBR policy with match ACL having the monitored applications and desired metric type for path determination. New/modified screens: New option in Interfaces page for enabling path monitoring: **Devices** > **Device Management** > **Edit Interfaces** > **Path Monitoring** > **Enable HTTP based Application Monitoring** check box.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Dual WAN/ISP Threat Defense Management Support	7.3.0	7.3.0	On a dual WAN-enabled threat defense, a single data interface was configured to communicate with the management center. Now, support to configure a secondary data interface is provided so that the communication channel is sustained when the primary data interface fails. The management center autoconfigures PBR to route the SF-Tunnel traffic from the *tapnlp* (internal) interface to one of the available data interfaces based on the priority and SLA metric.
Next-hop settings for PBR route map	7.3.0	7.1.0	You can configure the next-hops for the PBR route-map while enabling packet forwarding actions. New/modified screens: New fields in Add/Edit Forwarding Actions page for configuring egress interfaces: **Device Management** > **Routing** > **Policy Based Routing** > **Add Forwarding Actions** page.
PBR and Path Monitoring	7.2.0	7.2.0	PBR uses path monitoring to collect the performance metrics (RTT, jitter, packet-lost, and MOS) of the egress interfaces. You must enable path monitoring for the interface and configure the monitoring type. You can configure a PBR policy with the desired metric for path determination. New/modified screens: New tab in Interfaces page for enabling path monitoring: **Devices** > **Device Management** > **Edit Interfaces** > **Path Monitoring** tab.
Configure policy based routing from the FMC web interface.	7.1.0	7.1.0	**Upgrade impact. Redo FlexConfigs after upgrade.** You can now configure policy based routing (PBR) from the FMC web interface. This allows you to classify network traffic based on applications and to implement direct internet access (DIA) to send traffic to the internet from a branch deployment. You can define a PBR policy and configure it on ingress interfaces, specifying match criteria and egress interfaces. Network traffic that matches the access control policy is forwarded through the egress interface based on priority or the order as configured in the policy. This feature requires Version 7.1+ on both the FMC and the device. When you upgrade the FMC to Version 7.1+, existing policy based routing FlexConfigs are removed. After you upgrade your devices to Version 7.1+, redo your policy based routing configurations in the FMC web interface. For devices that you do not upgrade to Version 7.1+, redo the FlexConfigs and configure them to deploy "every time." New/modified screens: **Devices** > **Device Management** > **Routing** > **Policy Based Routing**

PART V

Objects and Certificates

- Object Management, on page 1351
- Certificates, on page 1485

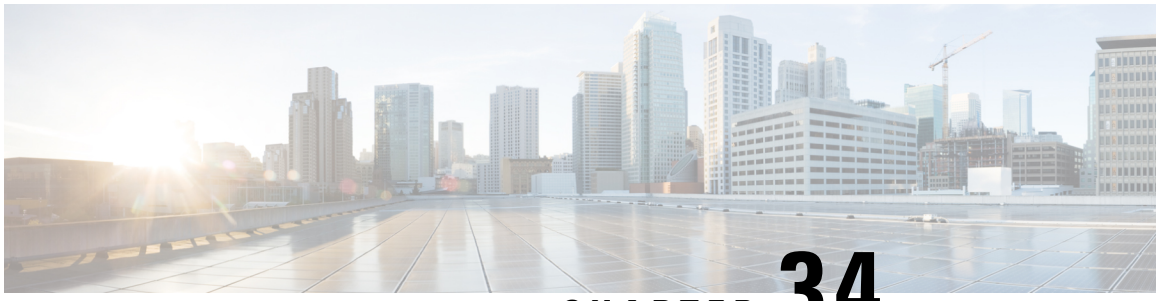

CHAPTER 34

Object Management

This chapter describes how to manage reusable objects.

- Introduction to Objects, on page 1352
- The Object Manager, on page 1354
- AAA Server, on page 1364
- Access List, on page 1369
- Address Pools, on page 1374
- Application Filters, on page 1376
- AS Path, on page 1376
- BFD Template, on page 1377
- Cipher Suite List, on page 1378
- Community List, on page 1379
- DHCP IPv6 Pool, on page 1381
- Distinguished Name, on page 1382
- DNS Server Group, on page 1384
- External Attributes, on page 1385
- File List, on page 1392
- FlexConfig, on page 1397
- Geolocation, on page 1398
- Interface, on page 1398
- Key Chain, on page 1399
- Network, on page 1401
- PKI, on page 1404
- Policy List, on page 1421
- Port, on page 1422
- Prefix List, on page 1424
- Route Map, on page 1426
- Security Intelligence, on page 1429
- Sinkhole, on page 1440
- SLA Monitor, on page 1441
- Time Range, on page 1442
- Time Zone, on page 1444
- Tunnel Zone, on page 1444
- URL, on page 1444

- Variable Set, on page 1446
- VLAN Tag, on page 1461
- VPN, on page 1462
- History for Object Management, on page 1480

Introduction to Objects

For increased flexibility and web interface ease-of-use, the system uses named *objects*, which are reusable configurations that associate a name with a value. When you want to use that value, use the named object instead. The system supports object use in various places in the web interface, including many policies and rules, event searches, reports, dashboards, and so on. The system provides many predefined objects that represent frequently used configurations.

Use the object manager to create and manage objects. Many configurations that use objects also allow you to create objects on the fly, as needed. You can also use the object manager to:

- View the policies, settings, and other objects where a network, port, VLAN, or URL object is used; see Viewing Objects and Their Usage, on page 1358.
- Group objects to reference multiple objects with a single configuration; see Object Groups, on page 1359.
- Override object values for selected devices; see Object Overrides, on page 1361.

After you edit an object used in an active policy, you must redeploy the changed configuration for your changes to take effect. You cannot delete an object that is in use by an active policy.

Note An object is configured on a managed device if, and only if, the object is used in a policy that is assigned to that device. If you remove an object from all policies assigned to a given device, the object is also removed from the device configuration on the next deployment, and subsequent changes to the object are not reflected in the device configuration.

Object Types

The following table lists the objects you can create in the system, and indicates whether each object type can be grouped or configured to allow overrides.

Object Type	Groupable?	Allows Overrides?
Network	yes	yes
Port	yes	yes
Interface: • Security Zone • Interface Group	no	no
Tunnel Zone	no	no
Application Filter	no	no

Object Type	Groupable?	Allows Overrides?
VLAN Tag	yes	yes
External Attribute: Security Group Tag (SGT) and Dynamic Object	no	no
URL	yes	yes
Geolocation	no	no
Time Range	no	no
Variable Set	no	no
Security Intelligence: Network, DNS, and URL lists and feeds	no	no
Sinkhole	no	no
File List	no	no
Cipher Suite List	no	no
Distinguished Name	yes	no
Public Key Infrastructure (PKI): • Internal and Trusted CA • Internal and External Certs	yes	no
Key Chain	no	yes
DNS Server Group	no	no
SLA Monitor	no	no
Prefix List: IPv4 and IPv6	no	yes
Route Map	no	yes
Access List: Standard and Extended	no	yes
AS Path	no	yes
Community List	no	yes
Policy List	no	yes
FlexConfig: Text and FlexConfig objects	no	yes

Objects and Multitenancy

In a multidomain deployment, you can create objects in Global and descendant domains with the exception of Security Group Tag (SGT) objects, which you can create only in the Global domain. The system displays

objects created in the current domain, which you can edit. It also displays objects created in ancestor domains, which you cannot edit, with the exception of security zones and interface groups.

Note Because security zones and interface groups are tied to device interfaces, which you configure at the leaf level, administrators in descendant domains can view and edit and groups created in ancestor domains. Subdomain users can add and delete interfaces from ancestor zones and groups, but cannot delete or rename the zones/groups.

Object names must be unique within the domain hierarchy. The system may identify a conflict with the name of an object you cannot view in your current domain.

For objects that support grouping, you can group objects in the current domain with objects inherited from ancestor domains.

Object overrides allow you to define device-specific or domain-specific values for certain types of object, including network, port, VLAN tag, and URL. In a multidomain deployment, you can define a default value for an object in an ancestor domain, but allow administrators in descendant domains to add override values for that object.

The Object Manager

You can use the object manager to create and manage objects and object groups.

The object manager displays 20 objects or groups per page. If you have more than 20 of any type of object or group, use the navigation links at the bottom of the page to view additional pages. You can also go to a specific page or click **Refresh** ( ) to refresh your view.

By default, the page lists objects and groups alphabetically by name. You can filter the objects on the page by name or value.

Importing Objects

Objects can be imported from a comma-separated values file. Up to 1000 objects can be imported in one attempt. The contents of the comma-separated values file should follow a specific format. The format is different for each object type. Only a few types of objects can be imported. See the following table to know the supported object types and the corresponding rules.

Object Type	Rules

Individual object	• The column header must be mentioned in capital letters.
	• The file must have the following columns headers:
	  • NAME
	  • DN
	• Both NAME and DN column entries are mandatory to import an entry.
	• You can import individual objects directly into an existing distinguished name object group.

Network object	• The column header must be mentioned in capital letters.
	• The file must have the following columns headers:
	  • NAME
	  • DESCRIPTION
	  • TYPE
	  • VALUE
	  • LOOKUP
	• The NAME and VALUE column entries are mandatory to import an entry of host, range, or network object type.
	• For an FQDN object, the TYPE column entry must mention 'fqdn,' and the LOOKUP column entry must be specified as 'ipv4,' 'ipv6,' or 'ipv4_ipv6.'
	• If no content is provided in the LOOKUP column entry for the FQDN object, then the object is saved with the ipv4_ipv6 field value.

Object Type	Rules

Port	• The column header must be mentioned in capital letters.
	• The file must have the following columns headers:
	  • NAME
	  • PROTOCOL
	  • PORT
	  • ICMPCODE
	  • ICMPTYPE
	• The NAME column entry is mandatory.
	• For 'tcp' and 'udp' protocol types, the PORT column entry is mandatory.
	• For 'icmp' and 'icmp6' protocol types, the ICMPCODE and ICMPTYPE column entries are mandatory.

URL	• The column header must be mentioned in capital letters.
	• The file must have the following columns headers:
	  • NAME
	  • DESCRIPTION
	  • URL
	• The NAME and URL column entries are mandatory to import an entry.

VLAN Tag	• The column header must be mentioned in capital letters.
	• The file must have the following columns headers:
	  • NAME
	  • DESCRIPTION
	  • TAG
	• The NAME and TAG column entries are mandatory to import an entry.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose one of the following object types from the left pane:

- **Distinguished Name** > **Individual Objects** >
- **Network Object**
- **Port**
- **URL**
- **VLAN Tag**

Step 3 Choose **Import Object** from the **Add [Object Type]** drop-down list.

Note
If you have selected **Individual Objects** in the previous step, click **Import**.

Step 4 Click **Browse**.

Step 5 Locate and select the comma-separated file on your system.

Step 6 Click **Open**.

Note
While importing **Distinguished Name** objects, you can optionally check the **Add imported Distinguished Name objects to the below object group** check box and select the group name from the drop-down box to import the objects directly to an existing distinguished name object group.

Step 7 Click **Import**.

Editing Objects

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose an object type from the list; see Introduction to Objects, on page 1352.

Step 3 Click **Edit** (✎) next to the object you want to edit.

If **View** (👁) appears instead, the object belongs to an ancestor domain and has been configured not to allow overrides, or you do not have permission to modify the object.

Step 4 Modify the object settings as desired.

Step 5 If you are editing a variable set, manage the variables in the set; see Managing Variables, on page 1458.

Step 6 For objects that can be configured to allow overrides:

- If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 1362. You can change this setting only for objects that belong to the current domain.
- If you want to add override values to this object, expand the Override section and click **Add**; see Adding Object Overrides, on page 1362.

Step 7 Click **Save**.

Step 8 If you are editing a variable set, and that set is in use by an access control policy, click **Yes** to confirm that you want to save your changes.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Viewing Objects and Their Usage

You can view usage details of objects on the Object Management page. Firewall Management Center provides this functionality for many object types. However, some object types are not supported.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose one of the following supported object types:

- Access List > Extended
- Access List > Standard
- AS Path
- Community List
- Interface
- Network
- Policy List
- Port
- Prefix List > IPv4 Prefix List
- Prefix List > IPv6 Prefix List
- Route Map
- SLA Monitor
- URL
- VLAN Tag

Step 3 Click the **Find Usage** ( ) icon next to the object.

The Object Usage window displays a list of all the policies, objects, and other settings where the object is in use. Click any of the listed items to know more about the object usage. For policies and some other settings where the object is used, you can click the corresponding links to visit the respective UI pages.

Filtering Objects or Object Groups

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Enter your filter criteria in the **Filter** field.

The page updates as you type to display matching items.

You can use the following wildcards:

- The asterisk (*) matches zero or more occurrences of a character.
- The caret (^) matches content at the beginning of a string.
- The dollar sign ($) matches content at the end of a string.

Step 3 Check the **Show Unused Object** check box to view the objects and the object groups that are unused anywhere in the system.

Note

- In case an object is a part of an unused object group, the object is considered as used. However, the unused object group is displayed when the **Show Unused Object** check box is checked.
- The **Show Unused Object** check box is available only for network, port, URL and VLAN tag object types.

Object Groups

Grouping objects allows you to reference multiple objects with a single configuration. The system allows you to use objects and object groups interchangeably in the web interface. For example, anywhere you would use a port object, you can also use a port object group.

You can group network, port, VLAN tag, URL, and PKI objects. Network object groups can be nested, that is, you can add a network object group to another network object group up to 10 levels.

Objects and object groups of the same type cannot have the same name.

When you edit an object group used in a policy (for example, a network object group used in an access control policy), you must re-deploy the changed configuration for your changes to take effect.

Deleting a group does not delete the objects in the group, just their association with each other. Additionally, you cannot delete a group that is in use in an active policy. For example, you cannot delete a VLAN tag group that you are using in a VLAN condition in a saved access control policy.

Grouping Reusable Objects

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 If the object type you want to group is **Network**, **Port**, **URL**, or **VLAN Tag**:

a) Choose the object type from the list of object types.
b) Choose **Add Group** from the **Add [Object Type]** drop-down list.

Step 3 If the object type you want to group is **Distinguished Name**:

a) Expand the **Distinguished Name** node.
b) Choose **Object Groups**.
c) Click **Add Distinguished Name Group**.

Step 4 If the object type you want to group is **PKI**:

a) Expand the **PKI** node.
b) Choose one of the following:

- **Internal CA Groups**
- **Trusted CA Groups**
- **Internal Cert Groups**
- **External Cert Groups**

c) Click **Add [Object Type] Group**.

Step 5 Enter a unique **Name**.

Step 6 Choose one or more objects from the list, and click **Add**.

You can also:

- Use the filter field **Search** ( ) to search for existing objects to include, which updates as you type to display matching items. Click **Reload** ( ) above the search field or click **Clear** ( ) in the search field to clear the search string.
- Click **Add** ( ) to create objects on the fly if no existing objects meet your needs.

Step 7 Optionally for **Network**, **Port**, **URL**, and **VLAN Tag** groups:

- Enter a **Description**.
- Check the **Allow Overrides** check box to allow overrides for this object group; see Allowing Object Overrides, on page 1362.

Step 8 Click **Save**.

What to do next

- If an active policy references your object group, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Object Overrides

An object override allows you to define an alternate value for an object, which the system uses for the devices you specify.

You can create an object whose definition works for most devices, and then use overrides to specify modifications to the object for the few devices that need different definitions. You can also create an object that needs to be overridden for all devices, but its use allows you to create a single policy for all devices. Object overrides allow you to create a smaller set of shared policies for use across devices without giving up the ability to alter policies when needed for individual devices.

For example, you might want to deny ICMP traffic to the different departments in your company, each of which is connected to a different network. You can do this by defining an access control policy with a rule that includes a network object called Departmental Network. By allowing overrides for this object, you can then create overrides on each relevant device that specifies the actual network where that device is connected.

You can target an object override to a specific domain. In this case, the system uses the object override value for all devices in the targeted domain unless you override it at the device level.

From the object manager, you can choose an object that can be overridden and define a list of device-level or domain-level overrides for that object.

You can use object overrides with the following object types only:

- Network
- Port
- VLAN tag
- URL
- SLA Monitor
- Prefix List
- Route Map
- Access List
- AS Path
- Community List
- Policy List
- Cert Enrollment (PKI)
- Key Chain

If you can override an object, the **Override** column appears for the object type in the object manager. Possible values for this column include:

- Green checkmark — indicates that you can create overrides for the object and no overrides have been added yet
- Red X — indicates that you cannot create overrides for the object
- Number — represents a count of the overrides that have been added to that object (for example, "2" indicates two overrides have been added)

Managing Object Overrides

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose from the list of object types; see Introduction to Objects, on page 1352.

Step 3 Click **Edit** (✎) next to the object you want to edit.

If **View** (👁) appears instead, the object belongs to an ancestor domain and has been configured not to allow overrides, or you do not have permission to modify the object.

Step 4 Manage the object overrides:

- Add—Add object overrides; see Adding Object Overrides, on page 1362.
- Allow—Allow object overrides; see Allowing Object Overrides, on page 1362.
- Delete—In the object editor, click **Delete** (🗑) next to the override you want to remove.
- Edit—Edit object overrides; see Editing Object Overrides, on page 1363.

Allowing Object Overrides

Procedure

Step 1 In the object editor, check the **Allow Overrides** check box.

Step 2 Click **Save**.

What to do next

Add object override values; see Adding Object Overrides, on page 1362.

Adding Object Overrides

Before you begin

Allow object overrides; see Allowing Object Overrides, on page 1362.

Procedure

Step 1 In the object editor, expand the **Override** section.

Step 2 Click **Add**.

Step 3 On **Targets**, choose domains or devices in the **Available Devices and Domains** list and click **Add**.

Step 4 On the Override tab, enter a **Name**.

Step 5 Optionally, enter a **Description**.

Step 6 Enter an override value.

Example:

For a network object, enter a network value.

Step 7 Click **Add**.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Editing Object Overrides

You can modify the description and the value of an existing override, but you cannot modify the existing target list. Instead, you must add a new override with new targets, which replaces the existing override.

Procedure

Step 1 In the object editor, expand the **Override** section.

Step 2 Click **Edit** (✎) next to the override you want to modify.

Step 3 Optionally, modify the **Description**.

Step 4 Modify the override value.

Step 5 Click **Save** to save the override.

Step 6 Click **Save** to save the object.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

AAA Server

Add reusable AAA server objects.

Add a RADIUS Server Group

RADIUS Server Group objects contain one or more references to RADIUS servers. These servers are used to authenticate users logging in through Remote Access VPN connections.

You can use this object with Firewall Threat Defense devices.

Before you begin

Note You cannot override RADIUS Server Group Objects.

Procedure

Step 1 Select **Objects** > **Object Management** > **AAA Server** > **RADIUS Server Group**.

All currently configured RADIUS Server Group objects will be listed. Use the filter to narrow down the list.

Step 2 Choose and edit a listed RADIUS Server Group object, or add a new one.

See RADIUS Server Options, on page 1366 and RADIUS Server Group Options, on page 1364 to configure this object.

Step 3 Click **Save**

RADIUS Server Group Options

Navigation Path

Objects > **Object Management** > **AAA Server** > **RADIUS Server Group**. Choose and edit a configured RADIUS Server Group object or add a new one.

Fields

- **Name** and **Description**—Enter a name and optionally, a description to identify this RADIUS Server Group object.
- **Group Accounting Mode**—The method for sending accounting messages to the RADIUS servers in the group. Choose **Single**, accounting messages are sent to a single server in the group, this is the default. Or, **Multiple**, accounting messages are sent to all servers in the group simultaneously.
- **Retry Interval**—The interval between attempts to contact the RADIUS servers. Values range from 1 to 10 seconds.

- **Realms**(Optional)—Specify or select the Active Directory (AD) realm this RADIUS server group is associated with. This realm is then selected in identity policies to access the associated RADIUS server group when determining the VPN authentication identity source for a traffic flow. This realm effectively provides a bridge from the identity policy to this Radius server group. If no realm is associated with this RADIUS server group, the RADIUS server group cannot be reached to determine the VPN authentication identity source for a traffic flow in an identity policy.

Note This field is mandatory if you use remote access VPN with User Identity and RADIUS as the identity source.

- **Enable authorize only**—If this RADIUS server group is not being used for authentication, but is being used for authorization or accounting, check this field to enable authorize-only mode for the RADIUS server group.

  Authorize only mode eliminates the need of including the RADIUS server password in the Access-Request. Thus, the password, configured for the individual RADIUS servers, is ignored.

- **Enable interim account update** and **Interval**—Enables the generation of RADIUS interim-accounting-update messages in order to inform the RADIUS server of newly assigned IP addresses. Set the length, in hours, of the interval between periodic accounting updates in the Interval field. The valid range is 1 to 120 and the default value is 24.

- **Enable Dynamic Authorization** and **Port**— Enables the RADIUS dynamic authorization or change of authorization (CoA) services for this RADIUS server group. Specify the listening port for RADIUS CoA requests in the **Port** field. The valid range is 1024 to 65535 and the default value is 1700. Once defined, the corresponding RADIUS server group will be registered for CoA notification and it listens to the port for the CoA policy updates from the Cisco Identity Services Engine (ISE).

- **Merge Downloadable ACL with Cisco AV Pair ACL**—Enables merging a downloadable access control list (dACL) with a Cisco attribute-value (AV) pair ACL.

  A downloadable ACL defines and updates access control lists in CiscoISE and allows ACL download to all the applicable controllers. For more information about using dACLs in Cisco ISE, see the chapter on Segmentation, section on authorization policies, in the *Cisco ISE Administrator Guide*.

  A Cisco AV pair ACL can be utilized to define specific authentication, authorization, and accounting elements for each individual session. For more information about using dACLs in Cisco ISE, see the chapter on Segmentation, section on authorization profile settings, in the *Cisco ISE Administrator Guide*.

  If you select **Merge Downloadable ACL with Cisco AV Pair ACL**, you can choose the following options:

  - **After Cisco AV Pair ACL** means the downloadable ACL entries should be placed after the Cisco AV pair entries.
  - **Before Cisco AV Pair ACL** means the downloadable ACL entries should be placed before the Cisco AV pair entries.

- **RADIUS Servers**—See RADIUS Server Options, on page 1366.

Related Topics

Add a RADIUS Server Group, on page 1364

RADIUS Server Options

Navigation Path

Objects > **Object Management** > **AAA Server** > **RADIUS Server Group** . Choose and edit a listed RADIUS Server Group object or add a new one. Then, in the RADIUS Server Group dialog, choose and edit a listed RADIUS Server or add a new one.

Fields

- **IP Address/Hostname**—The network object that identifies the hostname or IP address of the RADIUS server to which authentication requests will be sent. You may only select one, to add additional servers, add additional RADIUS Server to the RADIUS Server Group list.

Note The device now supports IPv6 IP addresses for RADIUS authentication.

- **RADIUS Server-Enabled Message Authenticator**—Message authenticators safeguard server-to-client communication and are required for a secure connection between your RADIUS server and firewall devices. Disabling message authenticators may expose your firewalls to potential attacks.

  Note the following:

  - Your RADIUS server must have the Message-Authenticator configuration.
  - You must have compatible Firewall Threat Defense devices that support message authenticators. Otherwise, your RADIUS logins may fail. For more information about the compatible FTDs that support this feature, see RADIUS Server-Enabled Message Authenticator Compatibility Matrix, on page 1367.

- **Authentication Port**—The port on which RADIUS authentication and authorization are performed. The default is 1812.
- **Key** and **Confirm Key**— The shared secret that is used to encrypt data between the managed device (client) and the RADIUS server.

  The key is a case-sensitive, alphanumeric string of up to 127 characters. Special characters are permitted.

  The key you define in this field must match the key on the RADIUS server. Enter the key again in the Confirm field.

- **Accounting Port**—The port on which RADIUS accounting is performed. The default is 1813.
- **Timeout**— Session timeout for authentication.

Note The timeout value must be 60 seconds or more for RADIUS two factor authentication. The default timeout value is 10 seconds.

- **Connect Using** —Establishes connectivity from the device to a RADIUS server using a route lookup or using a specific interface.
  - Click the **Routing** radio button to use the routing table.

- Click the **Specific Interface** radio button and choose a security zone/interface group or the Management interface (the default) from the drop-down list. If you want to use Management, you must choose it specifically; it is not available when using a route lookup. You cannot specify any other management-only interface as the RADIUS source. You can also choose a loopback interface group.

- **Redirect ACL**—Select the redirect ACL from the list or add a new one.

Note This is the name of the ACL defined in the device to decide the traffic to be redirected. The Redirect ACL name here must be the same as the *redirect-acl* name in ISE server. When you configure the ACL object, ensure that you select Block action for ISE and DNS servers, and Allow action for the rest of the servers.

Related Topics

Add a RADIUS Server Group, on page 1364
RADIUS Server Group Options, on page 1364

RADIUS Server-Enabled Message Authenticator Compatibility Matrix

Table 73: Firewall Management Center and Firewall Threat Defense Compatibility Matrix for RADIUS Server-Enabled Message Authenticator

Firewall Management Center Version	Firewall Threat Defense Version	NGIPS	FXOS Version	PAM RADIUS Version
7.7.10	• 7.2.10 or later • 7.6.1 or later • 7.7 or later	—	• 2.12.1.95 and later (Threat Defense 7.2.10 or later) • 2.16.0.136 and later (Threat Defense 7.6.1 or later) • 2.17.0.518 (Threat Defense 7.7 or later)	3.0.0

Add a Single Sign-on Server

Before you begin

Obtain the following from your SAML identity provider:

- Identity Provider Entity ID URL

- Sign-in URL
- Sign-out URL
- Identity provider certificate and enroll the certificate in Firewall Threat Defense using the Firewall Management Center web interface (**Devices** > **Certificates**)

For more information, see Configuring a SAML Single Sign-On Authentication, on page 1668.

Procedure

Step 1 Choose **Objects** > **Object Management** > **AAA Server** > **Single Sign-on Server**.

Step 2 Click **Add Single Sign-on Server** and provide the following details:

- **Name**—The name of the SAML single sign-on server object.
- **Identity Provider Entity ID**—The URL that is defined in SAML IdP to identify a service provider uniquely.

  The URL for a page that serves a metadata XML that describes how the SAML Issuer is going to respond to requests.
- **SSO URL**—The URL for signing into the SAML identity provider server.
- **Logout URL**—The URL for signing out of the SAML identity provider server.
- **Base URL**—URL that will redirect the user back to Firewall Threat Defense once the identity provider authentication is done. This is the URL of the access interface configured for the Firewall Threat Defense remote access VPN.
- **Identity Provider Certificate**—Certificate of the IdP enrolled into the Firewall Threat Defense to verify the messages signed by the IdP.

  Choose an identify provider certificate from the list or click **Add** to create a new certificate enrollment object.

  For more information, see Managing Firewall Threat Defense Certificates, on page 1486.

  You must enroll all of the Microsoft Azure registered application CA certificates as Trustpoints on the Firewall Threat Defense. The Microsoft Azure SAML identity provider is configured on Firewall Threat Defense for the initial application. All connection profiles are mapped to the configured MS Azure SAML identity provider. For each of the MS Azure applications (other than the default), you can choose the required trustpoint(CA certificate) in the connection profile configuration of the remote access VPN.

  For details, see Configure AAA Settings for Remote Access VPN, on page 1597.
- **Service Provider Certificate**—Firewall Threat Defense certificate, which will be used to sign the requests and build circle of trust with IdP.

  If you have not enrolled internal Firewall Threat Defense certificates, click + to add and enroll a certificate. For more information, see Managing Firewall Threat Defense Certificates, on page 1486.
- **Request Signature**—Select the encryption algorithm to sign the SAML single sign-on requests.

  The signatures are listed from weakest to strongest: SHA1,SHA256, SHA384, SHA512. Select None to disable encryption.

- **Request Timeout**—Specify the SAML assertion validity duration for the users to complete the single sign-on request. The SAML IdP has two time outs: *NotBefore* and *NotOnOrAfter*. The Firewall Threat Defense validates if its current time is within the time range of (lower limit) *NotBefore* and (upper limit) the smaller of *NotBefore* plus *timeout* and *NotOnOrAfter*. Thus, if you set a timeout longer than the IdP's *NotOnOrAfter* timeout, the specified timeout is ignored and the *NotOnOrAfter* timeout is selected. If the sum of the specified timeout and the *NotBefore* timeout is less than the *NotOnOrAfter* time, Firewall Threat Defense timeout overrides the timeout.

  The timeout range is 1-7200 seconds; the default is 300 seconds.

- **Enable IdP only accessible on Internal Network**—Select this option if the SAML IdP resides on the internal network. Firewall Threat Defense acts as a gateway and establishes communication between the users and IdP using an anonymous webvpn session.

- **Request IdP re-authentication on Login**—Select this option to authenticate user at each login even if the previous IdP session is valid.

- **Allow Overrides**—Select this check box to allow overrides for this single sign-on server object.

Step 3 Click **Save**.

Related Topics

Configure AAA Settings for Remote Access VPN, on page 1597

Access List

An access list object, also known as an access control list (ACL), selects the traffic to which a service will apply. You use these objects when configuring particular features, such as route maps, for Firewall Threat Defense devices. Traffic identified as allowed by the ACL is provided the service, whereas “blocked” traffic is excluded from the service. Excluding traffic from a service does not necessarily mean that it is dropped altogether.

You can configure the following types of ACL:

- Extended—Identifies traffic based on source and destination address and ports. Supports IPv4 and IPv6 addresses, which you can mix in a given rule.
- Standard—Identifies traffic based on destination address only. Supports IPv4 only.

An ACL is composed of one or more access control entry (ACE), or rule. The order of ACEs is important. When the ACL is evaluated to determine if a packet matches an “allowed” ACE, the packet is tested against each ACE in the order in which the entries are listed. After a match is found, no more ACEs are checked. For example, if you want to “allow” 10.100.10.1, but “block” the rest of 10.100.10.0/24, the allow entry must come before the block entry. In general, place more specific rules at the top of an ACL.

Packets that do not match an “allow” entry are considered to be blocked.

The following topics explain how to configure ACL objects.

Configure Extended ACL Objects

Use extended ACL objects when you want to match traffic based on source and destination addresses, protocol and port, application group or if the traffic is IPv6.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **Objects** > **Object Management** > **Access List** > **Extended** from the table of contents.

Step 2 Do one of the following:

- Click **Add Extended Access List** to create a new object.
- Click **Edit** (✎) to edit an existing object.

Step 3 In the New Extended Access List Object dialog box, enter a name for the object (no spaces allowed), and configure the access control entries:

a) Do one of the following:

- Click **Add** to create a new entry.
- Click **Edit** (✎) to edit an existing entry.

b) Select the **Action**, whether to Allow (match) or Block (not match) the traffic criteria.

Note
The **Logging**, **Log Level**, and **Log Interval** options are used for access rules only (ACLs attached to interfaces or applied globally). Because ACL objects are not used for access rules, leave these values at their defaults.

c) Configure the source and destination addresses on the **Network** tab using any of the following techniques:

- Select the desired network objects or groups from the Available list and click **Add to Source** or **Add to Destination**. You can create new objects by clicking the + button above the list. You can mix IPv4 and IPv6 addresses.
- Type an address in the edit box below the source or destination list and click **Add**. You can specify a single host address (such as 10.100.10.5 or 2001:DB8::0DB8:800:200C:417A), or a subnet (in 10.100.10.0/24 or 10.100.10.0 255.255.255.0 format, or for IPv6, 2001:DB8:0:CD30::/60).

d) Click the **Port** tab and configure the service using any of the following techniques.

- Select the desired port objects from the Available list and click **Add to Source** or **Add to Destination**. You can create new objects by clicking the + button above the list. The object can specify TCP/UDP ports, ICMP/ICMPv6 message types, or other protocols (including “any”). However, the source port, which you typically would leave empty, accepts TCP/UDP only. You cannot select port groups.

  For TCP/UDP, note that you must use the same protocol in both the source and destination fields, if you specify both. For example, you cannot specify a UDP source port and a TCP destination port.
- Type or select a port or protocol in the edit box below the source or destination list and click **Add**.

Note
To get an entry that applies to all IP traffic, select a destination port object that specifies “all” protocols.

e) Click the **Application** tab and choose the applications that are to be grouped for the direct internet access policy.

Important
- You cannot configure applications for cluster devices. Hence, this tab is not applicable for cluster devices.
- Use extended ACL with applications only in policy-based routing. Do not use it in other policies as its behavior is unknown and not supported. Ensure migration of the realm/ISE configuration for policy-based routing that uses User Identity and SGT in extended ACL.

Note
- The **Available Applications** list displays a fixed set of pre-defined applications. This list is a subset of the applications that are available on the Access Control policy as only they can be detected by their first packet (FQDN end-points resolved to IP addresses and port). The application definitions are updated through the VDB updates and are pushed to Firewall Threat Defense during subsequent deployments.
- User-defined custom applications or group of applications are not supported.
- Currently, Firewall Management Center neither supports user-defined custom applications or group of applications nor allows you to modify the pre-defined applications list.
- You can use the filter options provided under the **Application Filters** to refine this list.

f) Click the **Users** tab and choose the users, user groups, or both that are to be classified for the Policy Based Routing (PBR).

Important
Use extended ACL with users, user groups, or both only for Policy Based Routing. Do not use it in other policies as its behavior is unknown and not supported.

Note
- The **Available Realms** list displays the configured active directory/LDAP realms. For information on creation of realm and managing them, see Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426 and Manage a Realm, on page 2455 respectively.

  Note
  Local realms and Azure AD realms are not supported.

- The **Available Users** list displays the downloaded users and user groups of the selected AD/LDAP realms. To download the users, user groups, or both, navigate to **Integrations** > **Other Integrations** > **Realms**, and then click **Download** against the relevant active directory/LDAP realms.

  Note
  Threat defense can support a maximum of 512 user groups and 64000 user-IP mappings.

- The user to IP mapping and user group membership information are updated and pushed to the Firewall Threat Defense from the Firewall Management Center during the user login or logouts, and changes in the group memberships.

g) Click the **Security Group Tag** tab and choose the source SGT tags to be classified for the direct internet access policy.

Important
Use extended ACL with SGTs only for Policy Based Routing. Do not use it in other policies as its behavior is unknown and not supported.

Note
- The **Available Security Group Tags** list displays the configured security group tags. You can choose to use the ISE SGTs or create custom SGTs.
- To use ISE SGTs, ensure that ISE is integrated with Firewall Management Center with **Session Directory Topic** and SXP Topic subscribed. For more information on ISE integration, see Ways to Configure the Cisco Identity Services Engine (Cisco ISE) Identity Source, on page 2512.

  Note
  The supported ISE versions are 3.2, 3,1, 3.0, and 2.7 patch 2+.

- For information on creating custom SGTs, see Creating Security Group Tag Objects, on page 1392.

h) Select the required application, and click **Add to Rule**.

Note
- Do not configure destination networks and applications in the extended ACL object.
- The selected applications (Nertwork Service objects) in each of the access control entries, form a Network Service Group (NSG) and this group is deployed on the Firewall Threat Defense. The NSG is used in direct internet access to classify traffic based on the match with the selected application group.

i) Click **Add** to add the entry to the object.
j) If necessary, click and drag the entry to move it up or down in the rule order to the desired location.

Repeat the process to create or edit additional entries in the object.

Step 4 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 1362.

Step 5 Click **Save**.

Configure a Service Access Object

A Service Access object defines the conditions that traffic must meet to access a service, such as remote access VPN on the Firewall Threat Defense device. This object defines conditions as multiple rules to be executed in an order. Each service access rule has an Allow or Deny action and a set of match criteria such as country, continent, or user-defined geolocation objects. These rules manage access based on geolocations and ensure that only traffic from approved regions can access the specified services. If the traffic does not match any rules, the default action is enforced.

Before you begin

- Configure geolocation objects. For more information, see Geolocation, on page 1398.

- Configure a remote access VPN policy. For more information, see Create a New Remote Access VPN Policy, on page 1584.

Procedure

Step 1 Choose **Objects > Object Management**.

Step 2 In the left pane, click **Access List > Service Access**.

Step 3 Click **Add Service Access Object** to create a new object.

Step 4 In the **Add Service Access Object** dialog box, configure the following parameters:

a) In the **Name** field, enter a name for the object.

b) Click **Add Rule** to create service access rules for this object. In the **Add Service Access Rule** dialog box, configure the following parameters:

1. Choose an action as **Allow** or **Deny** from the drop-down list.
2. From **Available Countries**, select countries, continents, or user-defined geolocation objects and move them to the **Selected Geolocation** list.
3. Click **Add**.

After you create a service access rule, a sequence number is assigned to the rule. This number determines the execution order of the rule. You cannot reorder these rules.

Step 5 From the **Default Action** drop-down list, choose **Allow All Countries** or **Deny All Countries**.

Step 6 (Optional) Check the **Allow Overrides** check box and click + to configure overrides for the service access object for devices.

Note
When you define an override for a device or a domain, whenever the system configures this object on the device or domain, it uses the override value instead of the value defined in the original object.

Step 7 (Optional) In the **Add Service Access Override** dialog box, configure the following parameters:

a) From the **Target Devices and Domains** drop-down list, choose the required devices or domains for the override.

b) Verify if you require the existing rules. Otherwise, delete them.

c) Click **Add Rule** to create service access rules for this object override.

1. Choose an action as **Allow** or **Deny** from the drop-down list.
2. From **Available Countries**, select countries, continents, or user-defined geolocation objects and move them to the **Selected Geolocation** list.
3. Click **Add**.

d) From the **Default Action** drop-down list, choose **Allow All Countries** or **Deny All Countries**.

e) Click **Add**.

Step 8 Click **Save**.

What to do next

Configure the Service Access object in the remote access VPN policy. For more information, see Configure Access Interfaces for Remote Access VPN, on page 1613.

Configure Standard ACL Objects

Use standard ACL objects when you want to match traffic based on destination IPv4 address only. Otherwise, use extended ACLs.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **Objects** > **Object Management** > **Access List** > **Standard** from the table of contents.

Step 2 Do one of the following:

- Click **Add Standard Access List** to create a new object.
- Click **Edit** (✎) to edit an existing object.

Step 3 In the New Standard Access List Object dialog box, enter a name for the object (no spaces allowed), and configure the access control entries:

a) Do one of the following:

- Click **Add** to create a new entry.
- Click **Edit** (✎) to edit an existing entry.

b) For each access control entry, configure the following properties:

- **Action**—Whether to Allow (match) or Block (not match) the traffic criteria.
- **Network**—Add the IPv4 network objects or groups that identify the destination of the traffic.

c) Click **Add** to add the entry to the object.

d) If necessary, click and drag the entry to move it up or down in the rule order to the desired location.

Repeat the process to create or edit additional entries in the object.

Step 4 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 1362.

Step 5 Click **Save**.

Address Pools

You can configure IP address pools for both IPv4 and IPv6 that can be used for clustering or for VPN remote access profiles. For clustering in Individual interface mode, you can also configure MAC address pools.

Procedure

Step 1 Select **Objects** > **Object Management** > **Address Pools**.

Step 2 Click **IPv4 Pools** and then **Add IPv4 Pools**, and configure the following fields.

- **Name**—Enter the name of the address pool. It can be up to 64 characters
- **Description**—Add an optional description for this pool.
- **IP Address**—Enter a range of addresses available in the pool. Use dotted decimal notation and a dash between the beginning and the end address, for example: 10.10.147.100-10.10.147.177.
- **Mask**—Identifies the subnet on which this IP address pool resides.
- **Allow Overrides**—Check this check box to enable object overrides. Click the expand arrow to show the **Overrides** table. You can add a new override by clicking **Add**. See Object Overrides, on page 1361 for more information.

Step 3 Click **Save**.

Step 4 Click **IPv6 Pools** and then **Add IPv6 Pools**, and configure the following fields.

- **Name**—Enter the name of the address pool. It can be up to 64 characters
- **Description**—Add an optional description for this pool.
- **IPv6 Address**—Enter the first IP address available in the configured pool and the prefix length in bits. For example: 2001:DB8::1/64.
- **Number of Addresses**—Identifies the number of IPv6 addresses, starting at the Starting IP Address, that are in the pool.
- **Allow Overrides**—Check this check box to enable overrides. Click the expand arrow to show the **Overrides** table. You can add a new override by clicking **Add**. See Object Overrides, on page 1361 for more information.

Step 5 Click **Save**.

Step 6 Click **MAC Address Pool** and then **Add MAC Address Pool**, and configure the following fields.

For clustering in Individual interface mode, you can configure a MAC address pool for the interface. It is not common to manually configure MAC addresses for an interface, but if you have special needs to do so, then this pool is used to assign a unique MAC address to each interface. See Configure the MAC Address, on page 881.

- **Name**—Enter the name of the address pool. It can be up to 64 characters
- **Description**—Add an optional description for this pool.
- **MAC Address Range**—Enter a range of MAC addresses available in the pool. Use a dash between the beginning and the end address, for example: 000C.F142.4CD1-000C.F142.4CD7.

- **Allow Overrides**—Check this check box to enable overrides. Click the expand arrow to show the **Overrides** table. You can add a new override by clicking **Add**. See Object Overrides, on page 1361 for more information.

Application Filters

System-provided application filters help you perform application control by organizing applications according to basic characteristics: type, risk, business relevance, category, and tags. In the object manager, you can create and manage reuseable user-defined application filters based on combinations of the system-provided filters, or on custom combinations of applications. For detailed information, see Application Rule Conditions, on page 939.

AS Path

An AS Path is a mandatory attribute to set up BGP. It is a sequence of AS numbers through which a network can be accessed. An AS-PATH is a sequence of intermediate AS numbers between source and destination routers that form a directed route for packets to travel. Neighboring autonomous systems (ASes) use BGP to exchange and update messages about how to reach different AS prefixes. After each router makes a new local decision on the best route to a destination, it will send that route, or path information, along with the accompanying distance metrics and path attributes, to each of its peers. As this information travels through the network, each router along the path prepends its unique AS number to a list of ASes in the BGP message. This list is the route's AS-PATH. An AS-PATH along with an AS prefix, provides a specific handle for a one-way transit route through the network. Use the Configure AS Path page to create, copy and edit autonomous system (AS) path policy objects. You can create AS path objects to use when you are configuring route maps, policy maps, or BGP Neighbor Filtering. An AS path filter allows you to filter the routing update message by using regular expressions.

You can use this object with Firewall Threat Defense devices.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **AS Path** from the table of contents.

Step 2 Click **Add AS Path**.

Step 3 Enter a name for the AS Path object in the **Name** field. Valid values are between 1 and 500.

Step 4 Click **Add** on the **New AS Path Object** window.

a) Select the Allow or Block options from the **Action** drop-down list to indicate redistribution access.
b) Specify the regular expression that defines the AS path filter in the **Regular Expression** field.
c) Click **Add**.

Step 5 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 1362.

Step 6 Click **Save**.

BFD Template

The BFD template specifies a set of BFD interval values. BFD interval values as configured in the BFD template are not specific to a single interface. You can also configure authentication for single-hop and multi-hop sessions. Echo mode is disabled by default. You can enable Echo mode on single-hop only.

Procedure

Step 1 Choose **Objects** > **Object Management** > **BFD Template**.

Step 2 Click **Add BFD Template** or **Edit**.

Note
If you are editing a template, you cannot modify its name and type.

Step 3 On the **Template** tab, configure the following:

- **Template Name**—The name of this BFD template. You must assign a name in order to configure the rest of the parameters in the template. The template name cannot have spaces and cannot have only numbers.
- **Type**—Click the **Single-Hop** or **Multi-Hop** radio button.
- **Enable Echo**—(Optional) Enables Echo for the single-hop template.

If the Echo function is not negotiated, BFD control packets are sent at a high rate to meet the detection time. If the Echo function is negotiated, BFD control packets are sent at a slower, negotiated rate and self-directed echo packets are sent at a high rate. We recommend that you use Echo mode, if possible.

Step 4 On the **Interval** tab, configure the following:

a) From the **Interval Type** drop-down list, select **Microseconds** or **Milliseconds**.
b) In the **Multiplier** field, enter the value to be used for computing the hold down time. This value indicates the number of consecutive BFD control packets that must be missed from a BFD peer before BFD declares that the peer is unavailable and the Layer 3 BFD peer is informed of the failure. The range is 3 to 50. The default is 3.
c) In the **Minimum Transmit** field, enter the minimum transmit interval capability. The range is 50 to 999 milliseconds or 50000 to 999000 microseconds.
d) In the **Minimum Receive** field, enter the minimum receive interval capability. The range is 50 to 999 milliseconds or 50000 to 999000 microseconds.

Step 5 On the **Authentication** tab, configure the following:

- **Authentication Type**—Select **NONE**, **md5**, **meticulous-sha-1**, **meticulous-md5**, or **sha-1** from the drop-down list.
- **Encrypted Password**—(Optional) Enables encryption of the authentication password.
- **Password**—The authentication password that must be sent and received in the packets using the routing protocol being authenticated. The valid value is a string containing 1 to 29 uppercase and lowercase alphanumeric characters, except that the first character CANNOT be a digit or a digit followed by a whitespace. For example, '1password' or '0 password' is invalid.

- **Key ID**—The shared key ID that matches the key value. The range is 0 to 255.

Step 6 Click **OK**.

Step 7 Click **Apply** to save the BFD template configuration.

Cipher Suite List

A cipher suite list is an object comprised of several cipher suites. Each predefined cipher suite value represents a cipher suite used to negotiate an SSL- or TLS-encrypted session. You can use cipher suites and cipher suite lists in SSL rules to control encrypted traffic based on whether the client and server negotiated the SSL session using that cipher suite. If you add a cipher suite list to an SSL rule, SSL sessions negotiated with any of the cipher suites in the list match the rule.

Note Although you can use cipher suites in the web interface in the same places as cipher suite lists, you cannot add, modify, or delete cipher suites.

Creating Cipher Suite Lists

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Cipher Suite List** from the list of object types.

Step 3 Click **Add Cipher Suites**.

Step 4 Enter a **Name**.

Step 5 Choose one or more cipher suites from the **Available Ciphers** list.

Step 6 Click **Add**.

Step 7 Optionally, click **Delete** ( ) next to any cipher suites in the **Selected Ciphers** list that you want to remove.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Community List

A Community is an optional transitive BGP attribute. A community is a group of destinations that share some common attribute. It is used for route tagging. The BGP community attribute is a numerical value that can be assigned to a specific prefix and advertised to other neighbors. Communities can be used to mark a set of prefixes that share a common attribute. Upstream providers can use these markers to apply a common routing policy such as filtering or assigning a specific local preference or modifying other attributes. Use the Configure Community Lists page to create, copy and edit community list policy objects. You can create community list objects to use when you are configuring route maps or policy maps. You can use community lists to create groups of communities to use in a match clause of a route map. The community list is an ordered list of matching statements. Destinations are matched against the rules until a match is found.

You can use this object with Firewall Threat Defense devices.

Procedure

Step 1 Select **Objects** > **Object ManagementObjects** > **Object Management** and choose **Community List** from the table of contents.

Step 2 Click **Add Community List**.

Step 3 In the **Name** field, specify a name for the community list object.

Step 4 Click **Add** on the **New Community List Object** window.

Step 5 Select the **Standard** radio button to indicate the community rule type.

Standard community lists are used to specify well-known communities and community numbers.

Note
You cannot have entries using Standard and entries using Expanded community rule types in the same Community List object.

a) Select the Allow or Block options from the **Action** drop-down list to indicate redistribution access.
b) In the **Communities** field, specify a community number. Valid values can be from 1 to 4294967295 or from 0:1 to 65534:65535.
c) Select the appropriate **Route Type**.

- **Internet** — Select to specify the Internet well-known community. Routes with this community are advertised to all peers (internal and external).
- **No Advertise** — Select to specify the no-advertise well-known community. Routes with this community are not advertised to any peer (internal or external).
- **No Export** — Select to specify the no-export well-known community. Routes with this community are advertised to only peers in the same autonomous system or to only other sub-autonomous systems within a confederation. These routes are not advertised to external peers.

Step 6 Select the **Expanded** radio button to indicate the community rule type.

Expanded community lists are used to filter communities using a regular expression. Regular expressions are used to specify patterns to match COMMUNITIES attributes.

a) Select the Allow or Block options from the **Action** drop-down list to indicate redistribution access.
b) Specify the regular expression in the **Expressions** field.

Step 7 Click **Add**.

Step 8 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 1362.

Step 9 Click **Save**.

Extended Community

An extended community is a larger group of destinations that share some common attribute. The BGP extended community list has attributes that can be used to mark a set of prefixes that share a common attribute. These markers are used in the match clause of a route map to filter the routes for implementing route leaks among virtual routers. You can also define policy list objects with the extended community list for filtering. The extended community list is an ordered list of matching statements. Routes are matched against the rules until a match is found with the specified route target (standard) or regular expression (expanded). Use the Extended Community page to create and edit extended community list policy objects.

Note The extended community lists are applicable only for configuring import or export of routes.

You can use this object with Firewall Threat Defense devices.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **Community List** > **Extended Community** from the table of contents.

Step 2 Click **Add Extended Community List**.

Step 3 In the **Name** field, specify a name for the extended community list object. The length of the name cannot exceed 80 characters.

Step 4 Select the extended community rule type:

- Click the **Standard** radio button to specify one or more route targets.
- Click the **Expanded** radio button to specify regular expressions.

Note
You cannot have entries using Standard and Expanded extended community rule type in the same Extended Community List object.

Step 5 Click **Add**.

Step 6 If you have selected **Standard** as the extended community rule type, specify the following:

a) In the **Sequence No** field, enter the order in which you want the rule to be executed.

The sequence number must be unique in the list.

b) From the **Action** drop-down list, if you want to permit routes that have matching route target that is specified here, select **Allow**; if you want to deny routes that have matching route target that is specified here, select **Block**.

c) In the **Route Target** field, specify a route target.

- You can add a single route target or a set of route targets separated by commas in a single entry. For example, *1:2,1:4,1:6*.
- Valid values can be from 1:1 to 65534:65535.
- You can have a maximum of 8 route targets in an entry.
- You cannot have redundant route target set across multiple entries. For example, say you want to configure *seq1* with *1:200,100:100,1:300* route targets, and *seq2* with *1:300,100:100,1:200* route targets. This results in redundant route target set and cannot be deployed.

Step 7 If you have selected **Expanded** as the extended community rule type, specify the following:

a) In the **Sequence No** field, enter the order in which you want the rule to be executed.

The sequence number must be unique in the list.

b) From the **Action** drop-down list, if you want to permit routes that have matching regular expression that is specified here, select **Allow**; if you want to deny routes that have matching regular expression that is specified here, select **Block**.

c) Specify the regular expression in the **Expressions** field.

- You can add a single route target or a set of route targets separated by a space in a single entry. For example, *^((16) | (18)):(.)$*.
- You can add a maximum of 16 regular expressions to an entry.
- You cannot have redundant regular expression set across multiple entries. For example, say you want to configure *seq1* with *^((16) | (18)):(.)$ ^4_[0-9]*$* route targets, and *seq2* with *^4_[0-9]*$ ^((16) | (18)):(.)$* route targets. This results in redundant regular expression set and cannot be deployed.

For details on BGP regular expressions, see here.

Step 8 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 1362.

Step 9 Click **Save**.

The extended community list can be referenced in the match clause of the Route Map object or Policy List object:

- In the Route Map object, the name of the extended community list is displayed in the **Add Route Map Entry** > **Match Clause** > **BGP** > **Community List** > **Add Extended Community List** dialog. For more details on configuring BGP settings in a route map, see Route Map, on page 1426.
- In the Policy List object, the name of the extended community list is displayed in the **Add Policy List** > **Community Rule** > **Add Extended Community List** dialog. For more details on configuring BGP settings in a policy list, see Policy List, on page 1421.

DHCP IPv6 Pool

For clients that use StateLess Address Auto Configuration (SLAAC) in conjunction with the Prefix Delegation feature (Enable the IPv6 Prefix Delegation Client, on page 867), you can configure the Firewall Threat Defense

to provide information such as the DNS server or domain name when they send Information Request (IR) packets to the Firewall Threat Defense by defining a DHCP IPv6 Pool and assigning it to the DHCPv6 server. The Firewall Threat Defense only accepts IR packets and does not assign addresses to the clients. You will configure the client to generate its own IPv6 address by enabling IPv6 autoconfiguration on the client. Enabling stateless autoconfiguration on a client configures IPv6 addresses based on prefixes received in Router Advertisement messages; in other words, based on the prefix that the Firewall Threat Defense received using Prefix Delegation.

To add a pool, see Create the DHCP IPv6 Pool, on page 914.

Distinguished Name

Each distinguished name object represents the distinguished name for a public key certificate's subject or issuer. You can use distinguished name objects and groups in TLS/SSL rules to control encrypted traffic based on whether the client and server negotiated the TLS/SSL session using a server certificate with the distinguished name as subject or issuer.

(A *distinguished name group* is a named collection of existing distinguished name objects.)

The distinguished name can consist of country code, common name, organization, and organizational unit, but typically consists of a common name only. For example, the common name in the certificate for `https://www.cisco.com` is `cisco.com`. (However, it's not always this simple; Distinguished Name (DN) Rule Conditions, on page 2323 shows how to find common names.) The certificate can contain multiple Subject Alternative Names (SANs) you can use as DNs in a rule condition. For detailed information about SANs, see RFC 5280, section 4.2.1.6.

The format of a distinguished name object that references a common name is `CN=`*`name`*. If you add a DN rule condition without `CN=`, the system prepends `CN=` before saving the object.

As discussed further in Distinguished Name (DN) Rule Conditions, on page 2323, the system uses Server Name Indication (SNI) to match the DN in the TLS/SSL rule whenever possible.

You can also add a distinguished name with one of each of the attributes listed in the following table, separated by commas.

Table 74: Distinguished name attributes

Attribute	Description	Allowed Values
C	Country Code	two alphabetic characters
CN	Common Name	up to 64 alphanumeric, backslash (/), hyphen (-), quotation ("), or asterisk (*) characters, or spaces
O	Organization	up to 64 alphanumeric, backslash (/), hyphen (-), quotation ("), or asterisk (*) characters, or spaces
OU	Organizational Unit	up to 64 alphanumeric, backslash (/), hyphen (-), quotation ("), or asterisk (*) characters, or spaces

Important notes about DN rule conditions

- The first time the system detects an encrypted session to a new server, DN data is not available for ClientHello processing, which *might* result in an undecrypted first session.

If the server requests TLS 1.3, the setting for TLS server identity discovery can help by making sure the server certificate is known before making decryption policy decisions. For more information, see Access Control Policy Advanced Settings, on page 1756.

- You *cannot* configure a distinguished name condition if you also choose the **Decrypt - Known Key** action. Because that action requires you to choose a server certificate to decrypt traffic, the certificate already matches the traffic.

Wildcard examples

You can define one or more asterisks (*****) as wildcards in an attribute. In a common name attribute, you can define one or more asterisks per domain name label. wildcards match only in that label, but you can define multiple labels with wildcards. See the following table for examples.

Table 75: Common Name attribute wildcard examples

Attribute	Matches	Does Not Match
CN=*ample.com	example.com	mail.example.com example.text.com ampleexam.com
CN=exam*.com	example.com	mail.example.com example.text.com ampleexam.com
CN=*xamp*.com	example.com	mail.example.com example.text.com ampleexam.com
CN=*.example.com	mail.example.com	www.myhost.example.com example.com example.text.com ampleexam.com

Note The DN object `CN=amp.cisco.com` would *not* match a CN like `CN=auth.amp.cisco.com`, which is why we recommend wildcards in these cases.

For more information and examples, see Distinguished Name (DN) Rule Conditions, on page 2323.

Related Topics

Distinguished Name (DN) Rule Conditions, on page 2323

Creating Distinguished Name Objects

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **Distinguished Name** node, and choose **Individual Objects**.

Step 3 Click **Add Distinguished Name**.

Step 4 Enter a **Name**.

Step 5 In the **DN** field, enter a value for the distinguished name or common name. You have the following options:

- If you add a distinguished name, you can include one of each attribute listed in Distinguished Name, on page 1382 separated by commas.
- If you add a common name, you can include multiple labels and wild cards.

Step 6 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

DNS Server Group

Domain Name System (DNS) servers resolve fully-qualified domain names (FQDN), such as www.example.com, to IP addresses.

Creating DNS Server Group Objects

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Click **DNS Server Group** from the network objects list.

Step 3 Click **Add DNS Server Group**.

Step 4 Enter a **Name**.

Step 5 Optionally, enter the **Default Domain** that will be used to append to the host names that are not fully-qualified.

This setting is only used for the default server group.

Step 6 The default **Timeout** and **Retries** values are pre-populated. Change these values if necessary.

- Retries—The number of times, from 0 to 10, to retry the list of DNS servers when the system does not receive a response. The default is 2.

- Timeout—The number of seconds, from 1 to 30, to wait before trying the next DNS server. The default is 2 seconds. Each time the system retries the list of servers, this timeout doubles.

Step 7 Enter the **DNS Servers** that will be a part of this group, either in IPv4 or IPv6 format as comma separated entries.

A maximum of 6 DNS servers can belong to one group.

Step 8 Click **Save**.

What to do next

The DNS servers configured in the DNS server group should be assigned to interface objects in the DNS platform settings. For more information, see DNS, on page 949.

External Attributes

Dynamic Objects

A *dynamic object* is an object that specifies one or many IP addresses retrieved either using REST API calls or using the Cisco Secure Dynamic Attributes Connector, which is capable of updating IP addresses from cloud sources. These dynamic objects can be used in access control rules without the need to deploy dynamic changes to the objects.

Note Unlike most other objects, dynamic objects do not have to be deployed to managed devices to take effect. Just add dynamic objects to the **Dyamic Attributes** tab of the rule and then deploy the rule; the object values are automatically updated on the managed device as soon as possible after being pushed by the Cisco Secure Dynamic Attributes Connector.

There are the following kinds of dynamic objects:

- Dynamic objects created using the dynamic attributes connector are pushed to the Firewall Management Center as soon as they're created and are updated at a regular interval.
- API-created dynamic objects:
  - Are IP addresses, with or without or classless inter-domain routing (CIDR), that can be used in access control rules much like a network object.
  - Do not support fully-qualified domain names or address ranges.
  - Must be updated using an API.

For more information about API-created dynamic objects, see About API-Created Dynamic Objects, on page 1390.

Create Dynamic Objects for the First Time

The page available at **Objects** > **Object Management** > **External Attributes** > **Dynamic Object** is displayed as follows if you have not configured any dynamic objects yet.

If you have already created some dynamic objects, see Work With Dynamic Objects, on page 1390.

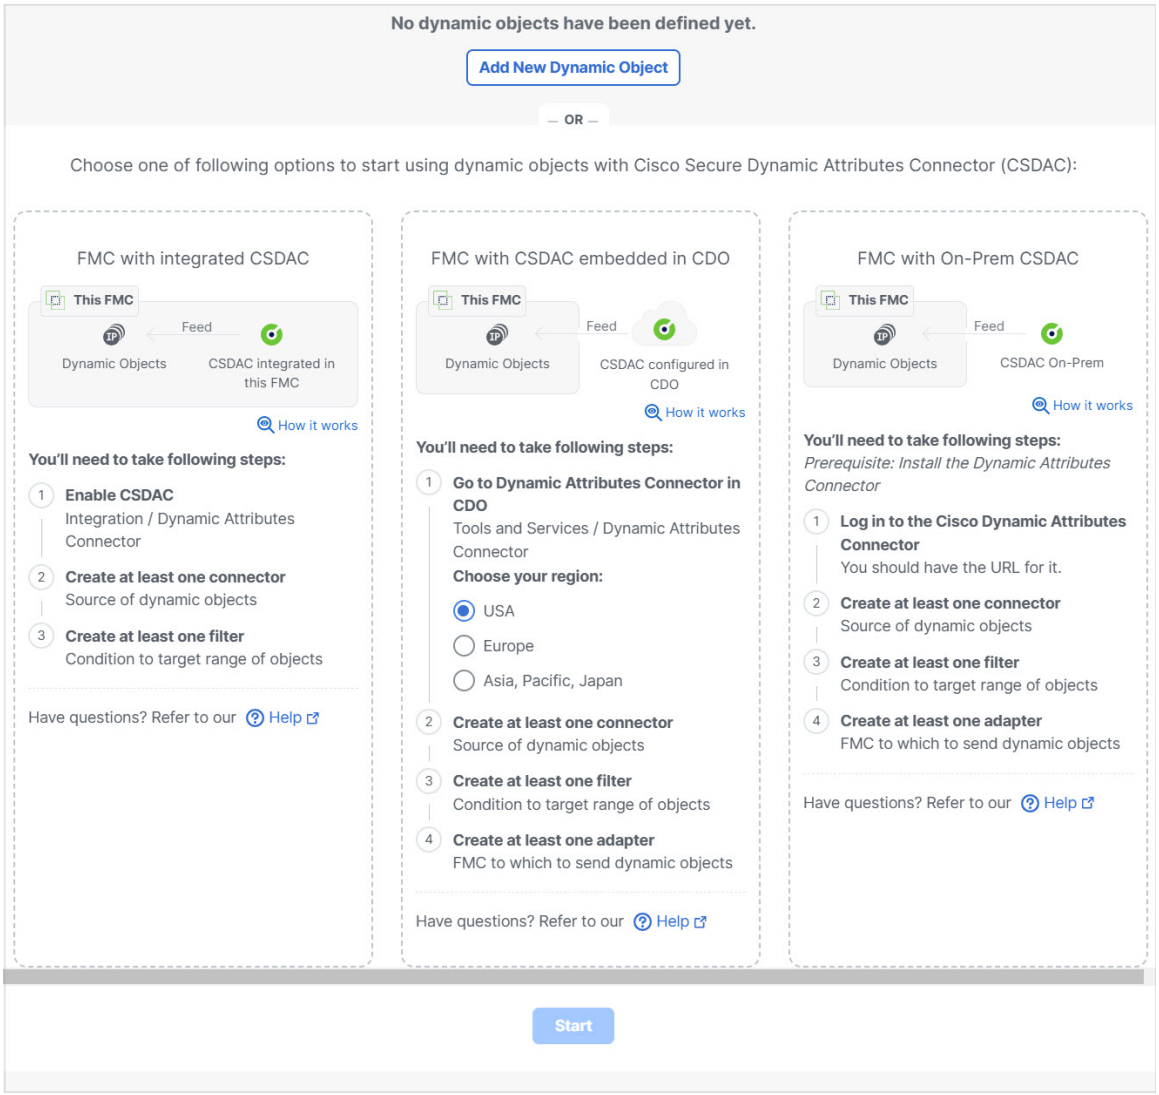

To use this page:

- Click the section that represents how you plan to deploy the dynamic attributes connector.
- Click **How it works** to see a diagram and get more help with that type of deployment.

  The following topics provide more information.

  - Create Dynamic Objects with the Embedded Cisco Secure Dynamic Attributes Connector, on page 1387
  - Create Dynamic Objects with Cisco Security Cloud Control, on page 1388

- Create Dynamic Objects with the On-Premises Cisco Secure Dynamic Attributes Connector, on page 1389

- Click **Start** to open the relevant application as follows:
  - If you clicked the **FMC with Integrated CSDAC** method, **Start** takes you to **Objects** > **Object Management** > **External Attributes** > **Dynamic Object** where you can enable and configure the Cisco Secure Dynamic Attributes Connector.

    For more information, see Enable the Cisco Secure Dynamic Attributes Connector, on page 1806.
  - If you clicked the **FMC with CSDAC embedded in** Security Cloud Control method, **Start** starts Cisco Security Cloud Control.

    For more information, see the guide titled *Managing Firewall Threat Defense with Cloud-delivered Firewall Management Center in Security Cloud Control Firewall Management* .
  - If you clicked the **FMC with On-Prem CSDAC** method, **Start** enables you to download the Cisco Secure Dynamic Attributes Connector if you have not done so already.

    For more information, see the Cisco Secure Dynamic Attributes Connector Configuration Guide.

Create Dynamic Objects with the Embedded Cisco Secure Dynamic Attributes Connector

The following page is displayed if you indicated you're configuring the Cisco Secure Dynamic Attributes Connector provided with this Secure Firewall Management Center. This Secure Firewall Management Center already has the Cisco Secure Dynamic Attributes Connector integrated with it (**Integration** > **Dynamic Attributes Connector**).

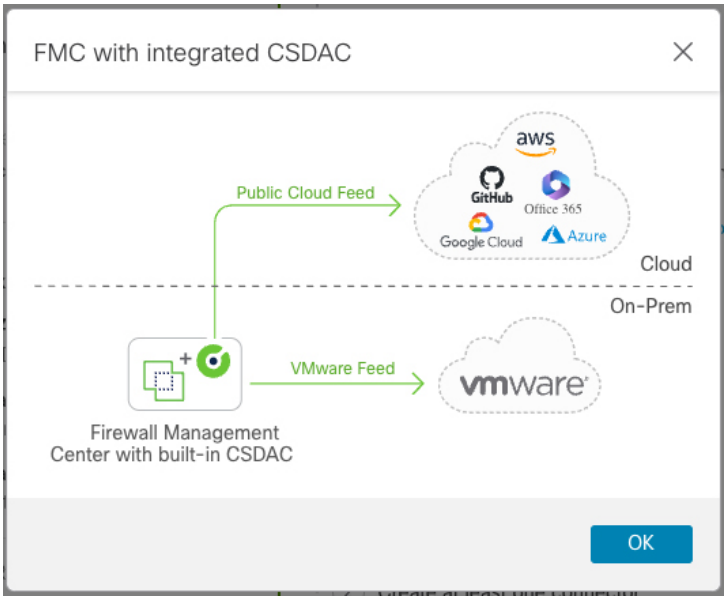

To use this type of deployment:

1. Enable the Cisco Secure Dynamic Attributes Connector as discussed in Enable the Cisco Secure Dynamic Attributes Connector, on page 1806.
2. Configure connectors, which retrieve IP addresses from cloud services.

For more information, see Create a Connector, on page 1816.

3. Configure dynamic attributes filters, which determine what IP addresses to send to the management center.

   For more information, see Create Dynamic Attributes Filters, on page 1839.

4. View your dynamic objects at **Objects** > **Object Management** > **External Attributes** > **Dynamic Object**.

5. Use dynamic objects in access control rules (**Policies** > **Access Control heading** > **Access Control**, then click the **Dynamic Attributes** tab).

Create Dynamic Objects with Cisco Security Cloud Control

The following page is displayed if you indicated you're configuring the Cisco Secure Dynamic Attributes Connector provided with Cisco Security Cloud Control.

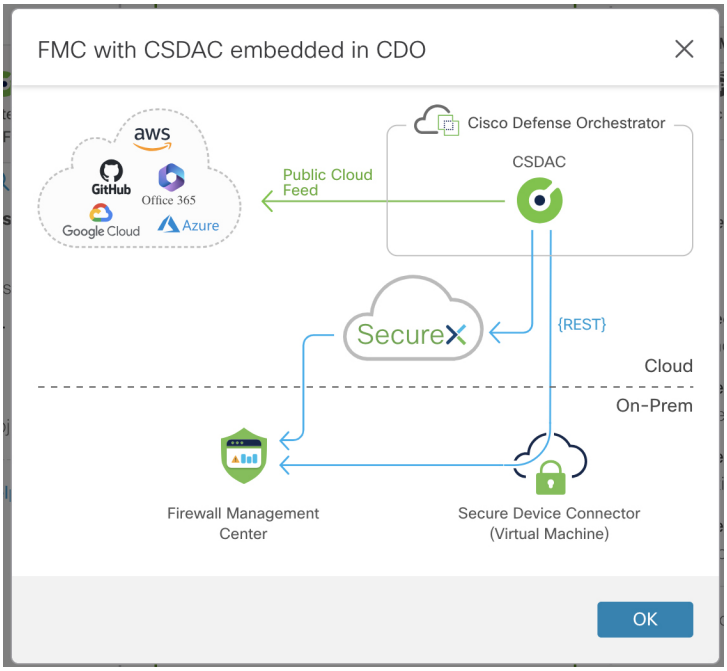

The preceding diagram has details about configuring Cisco Security Cloud Control that are not discussed in this guide. For more detailed information, see Secure Device Connector (SDC) or SecureX and Security Cloud Control.

To use this type of deployment:

1. Configure connectors, which retrieve IP addresses from cloud services.

   For more information, see Create a Connector, on page 1816.

2. Configure dynamic attributes filters, which determine what IP addresses to send to the management center.

   For more information, see Create Dynamic Attributes Filters, on page 1839.

3. Configure adapters, which send IP addresses to a Secure Firewall Management Center or Cloud-Delivered Firewall Management Center.

For more information, see the section on creating adapters in the Managing Firewall Threat Defense with Cloud-Delivered Firewall Management Center in Cisco Security Cloud Control.

4. Log in to the Secure Firewall Management Center you defined as an adapter.

   If the Secure Firewall Management Center is managed by Cisco Security Cloud Control, click and choose **Cloud-Delivered FMC**.

5. View your dynamic objects at **Objects** > **Object Management** > **External Attributes** > **Dynamic Object**.

6. Use dynamic objects in access control rules (**Policies** > **Access Control heading** > **Access Control**, then click the **Dynamic Attributes** tab).

Create Dynamic Objects with the On-Premises Cisco Secure Dynamic Attributes Connector

The following page is displayed if you indicated you're configuring the on-premises Cisco Secure Dynamic Attributes Connector to send dynamic objects to a Secure Firewall Management Center or Cloud-Delivered Firewall Management Center.

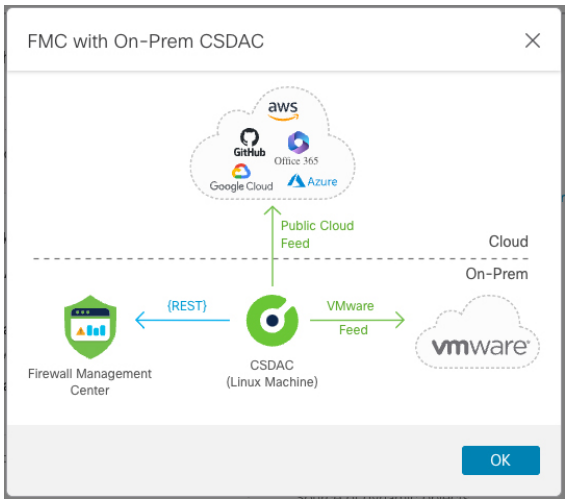

To use this type of deployment:

1. Install the Cisco Secure Dynamic Attributes Connector on a supported Linux virtual machine.

2. Configure connectors, which retrieve IP addresses from cloud services.

   For more information, see the section on creating connectors in the Cisco Secure Dynamic Attributes Connector Configuration Guide.

3. Configure adapters, which send IP addresses to a Secure Firewall Management Center or Cloud-Delivered Firewall Management Center.

   For more information, see the section on creating adapters in the Cisco Secure Dynamic Attributes Connector Configuration Guide.

4. Configure dynamic attributes filters, which determine what IP addresses to send to the management center.

   For more information, see the section on configuring dynamic attributes filters in the Cisco Secure Dynamic Attributes Connector Configuration Guide.

5. View your dynamic objects at **Objects** > **Object Management** > **External Attributes** > **Dynamic Object**.
6. Use dynamic objects in access control rules (**Policies** > **Access Control heading** > **Access Control**, then click the **Dynamic Attributes** tab).

For more information, see the Cisco Secure Dynamic Attributes Connector Configuration Guide.

Work With Dynamic Objects

The page available at **Objects** > **Object Management** > **External Attributes** > **Dynamic Object** is displayed similar to the following if you have already configured some dynamic objects.

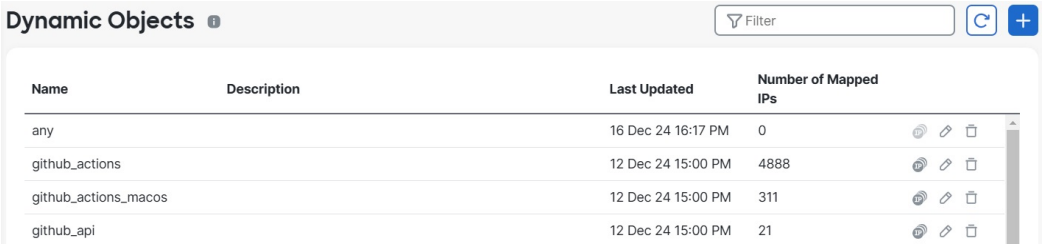

This page displays information about each dynamic object and enables you to view or download IP addresses associated with that object. For more information, see Dynamic Object Mappings, on page 1390.

Dynamic Object Mappings

If you configured dynamic objects either using the API or using the dynamic attributes connector, your connectors send IPs matching dynamic attributes filters to the Firewall Management Center at regular intervals.

To view or download a current list of these IP addresses, click **Show Mapped IDs** as the following figure shows.

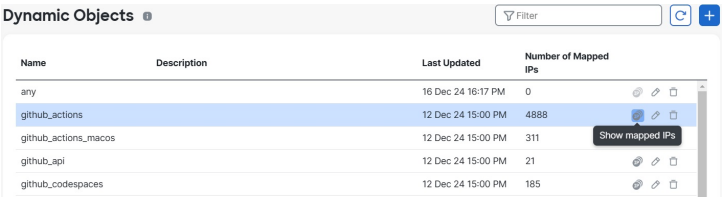

IP addresses are added dynamically with time so you should consider doing this on a regular basis, especially if your access control rules are not behaving as expected.

Related topics

- About API-Created Dynamic Objects, on page 1390
- About the Cisco Secure Dynamic Attributes Connector, on page 1803

About API-Created Dynamic Objects

A *dynamic object* is an object that specifies one or many IP addresses retrieved either using REST API calls or using the Cisco Secure Dynamic Attributes Connector, which is capable of updating IP addresses from cloud sources. These dynamic objects can be used in access control rules without the need to deploy dynamic changes to the objects.

For more information about the dynamic attributes connector, see the *Cisco Secure Dynamic Attributes Configuration Guide* (link to guide).

Differences between dynamic objects and network objects follow:

- Dynamic objects created using the dynamic attributes connector are pushed to the Firewall Management Center as soon as they're created and are updated at a regular interval.
- API-created dynamic objects:
  - Are IP addresses, with or without or classless inter-domain routing (CIDR), that can be used in access control rules much like a network object.
  - Do not support fully-qualified domain names or address ranges.
  - Must be updated using an API.

Related Topics

Add or Edit an API-Created Dynamic Object, on page 1391

Add or Edit an API-Created Dynamic Object

This procedure discusses how to add or edit a *dynamic object,* which is a group of IP addresses using the API, with or without or classless inter-domain routing (*CIDR*), that can be used in access control rules much like a network object.

Note This procedure is not necessary if you use the Cisco Secure Dynamic Attributes Connector because it automatically creates dynamic objects for you.

Before you begin

Consult the *Firepower Management Center REST API Quick Start Guide* for information about using the object services REST API to populate the IP object with an address. Dynamic objects do not require deployment.

Procedure

Step 1 Click **Objects** > **Object Management**.

Step 2 Click **External Attributes** > **Dynamic Objects**.

Step 3 Click **Add Dynamic Object** or **Edit** ( ).

Step 4 Enter a **Name** for the object and an optional **Description**.

Step 5 From the **Type** list, click **IP**.

What to do next

If necessary, update the dynamic object using the API. Deployment is not required.

Security Group Tag

A Security Group Tag (SGT) object specifies a single SGT value. You can use SGT objects in rules to control traffic with SGT attributes that were **not** assigned by Cisco ISE. You cannot group or override SGT objects.

Related Topics

Autotransition from Custom SGTs to ISE SGTs
Custom SGT Conditions
ISE SGT vs Custom SGT Rule Conditions

Creating Security Group Tag Objects

You can create these objects in the global domain only. To use the object on Classic devices, you must have the Control license. For Smart Licensed devices, any license will do.

Before you begin

- Disable ISE/ISE-PIC connections. You cannot create custom SGT objects if you use ISE/ISE-PIC as an identity source.

Procedure

Step 1 Click **Objects** > **Object Management**.
Step 2 Click **External Attributes** > **Security Group Tag**.
Step 3 Click **Add Security Group Tag**.
Step 4 Enter a **Name**.
Step 5 Optionally, enter a **Description**.
Step 6 In the **Tag** field, enter a single SGT.
Step 7 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

File List

If you use malware defense, and the AMP cloud incorrectly identifies a file's disposition, you can add the file to a *file list* to better detect the file in the future. These files are specified using SHA-256 hash values. Each file list can contain up to 10000 unique SHA-256 values.

There are two predefined categories of file lists:

Clean List

If you add a file to this list, the system treats it as if the AMP cloud assigned a clean disposition.

Custom Detection List

If you add a file to this list, the system treats it as if the AMP cloud assigned a malware disposition.

Because you manually specify the blocking behavior for the files included in these lists, the system does not query the AMP cloud for these files' dispositions. You must configure a rule in the file policy with either a **Malware Cloud Lookup** or **Block Malware** action and a matching file type to calculate a file's SHA value.

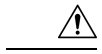

Caution Do **not** include malware on the clean list. The clean list overrides both the AMP cloud and the custom detection list.

Source Files for File Lists

You can add multiple SHA-256 values to a file list by uploading a comma-separated value (CSV) source file containing a list of SHA-256 values and descriptions. The Firewall Management Center validates the contents and populates the file list with valid SHA-256 values.

The source file must be a simple text file with a .csv file name extension. Any header must start with a pound sign (`#`); it is treated as a comment and not uploaded. Each entry should contain a single SHA-256 value followed by a description and end with either the `LF` or `CR+LF` Newline character. The system ignores any additional information in the entry.

Note the following:

- Deleting a source file from the file list also removes all associated SHA-256 hashes from the file list.
- You cannot upload multiple files to a file list if the successful source file upload results in the file list containing more than 10000 distinct SHA-256 values.
- The system truncates descriptions exceeding 256 characters to the first 256 characters on upload. If the description contains commas, you must use an escape character (`\,`). If no description is included, the source file name is used instead.
- All non-duplicate SHA-256 values are added to the file list. If a file list contains a SHA-256 value, and you upload a source file containing that value, the newly uploaded value does not modify the existing SHA-256 value. When viewing captured files, file events, or malware events related to the SHA-256 value, any threat name or description is derived from the individual SHA-256 value.
- The system does not upload invalid SHA-256 values in a source file.
- If multiple uploaded source files contain an entry for the same SHA-256 value, the system uses the most recent value.
- If a source file contains multiple entries for the same SHA-256 value, the system uses the last one.
- You cannot directly edit a source file within the object manager. To make changes, you must first modify your source file directly, delete the copy on the system, then upload the modified source file.
- The number of entries associated with a source file refers to the number of distinct SHA-256 values. If you delete a source file from a file list, the total number of SHA-256 entries the file list contains decreases by the number of valid entries in the source file.

Adding Individual SHA-256 Values to File Lists

You must have the Malware Defense license for this procedure.

You can submit a file's SHA-256 value to add it to a file list. You cannot add duplicate SHA-256 values.

Before you begin

- Right-click a file or malware event from the event view, choose **Show Full Text** in the context menu, and copy the full SHA-256 value for pasting into the file list.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **File List** from the list of object types.

Step 3 Click **Edit** (✎) next to the clean list or custom detection list where you want to add a file.

If **View** (👁) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 Choose `Enter SHA Value` from the **Add by** drop-down list.

Step 5 Enter a description of the source file in the **Description** field.

Step 6 Enter or paste the file's entire value in the **SHA-256** field. The system does not support matching partial values.

Step 7 Click **Add**.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Note After configuration changes are deployed, the system no longer queries the AMP cloud for files on the list.

Uploading Individual Files to File Lists

You must have the Malware Defense license for this procedure.

If you have a copy of the file you want to add to a file list, you can upload the file to the Secure Firewall Management Center for analysis; the system calculates the file's SHA-256 value and adds the file to the list. The system does not enforce a limit on the size of files for SHA-256 calculation.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **File List** from the list of object types.

Step 3 Click **Edit** (✎) next to the clean list or custom detection list where you want to add a file.

If **View** (👁) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 From the **Add by** drop-down list, choose **Calculate SHA**.

Step 5 Optionally, enter a description of the file in the **Description** field. If you do not enter a description, the file name is used for the description on upload.

Step 6 Click **Browse**, and choose a file to upload.

Step 7 Click **Calculate and Add SHA**.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Note After you deploy configuration changes, the system no longer queries the AMP cloud for files on the list.

Uploading Source Files to File Lists

You must have the Malware Defense license for this procedure.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Click **File List**.

Step 3 Click **Edit** (✎) next to the file list where you want to add values from a source file.

If **View** (👁) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 In the **Add by** drop-down list, choose `List of SHAs`.

Step 5 Optionally, enter a description of the source file in the **Description** field. If you do not enter a description, the system uses the file name.

Step 6 Click **Browse** to browse to the source file, then click **Upload and Add List**.

Step 7 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Note After you deploy the policies, the system no longer queries the AMP cloud for files on the list.

Editing SHA-256 Values in File Lists

You must have the Malware Defense license for this procedure.

You can edit or delete individual SHA-256 values on a file list. Note that you cannot directly edit a source file within the object manager. To make changes, you must first modify your source file directly, delete the copy on the system, then upload the modified source file.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Click **File List**.

Step 3 Click **Edit** (✎) next to the clean list or custom detection list where you want to modify a file.

If **View** (👁) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 You can:

- Click **Edit** (✎) next to the SHA-256 value you want to change, and modify the **SHA-256** or **Description** values as desired.
- Click **Delete** (🗑) next to the SHA-256 value you want to delete.

Step 5 Click **Save** to update the file entry in the list.

Step 6 Click **Save** to save the file list.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Note After configuration changes are deployed, the system no longer queries the AMP cloud for files on the list.

Downloading Source Files from File Lists

You must have the Malware Defense license for this procedure.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **File List** from the list of object types.

Step 3 Click **Edit** (✎) next to the clean list or custom detection list where you want to download a source file.

If **View** (👁) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 Next to the source file you want to download, click **View** (👁).

Step 5 Click **Download SHA List** and follow the prompts to save the source file.

Step 6 Click **Close**.

FlexConfig

Use FlexConfig policy objects in FlexConfig policies to provide customized configuration of features on Firewall Threat Defense devices that you cannot otherwise configure using Secure Firewall Management Center. For more information on FlexConfig policies, see FlexConfig Policy Overview, on page 2667.

You can configure the following types of objects for FlexConfig.

Text Objects

Text objects define free-form text strings that you use as variables in a FlexConfig object. These objects can have single values or be a list of multiple values.

There are several predefined text objects that are used in the predefined FlexConfig objects. If you use the associated FlexConfig object, you simply need to edit the contents of the text object to customize how the FlexConfig object configures a given device. When editing a predefined object, it is in general a better option to create device overrides for each device you are configuring, rather than directly change the default values of these objects. This helps avoid unintended consequences if another user wants to use the same FlexConfig object for a different set of devices.

For information on configuring text objects, see Configure FlexConfig Text Objects, on page 2694.

FlexConfig Objects

FlexConfig Objects include device configuration commands, variables, and scripting language instructions. During configuration deployment, these instructions are processed to create a sequence of configuration commands with customized parameters to configure specific features on the target devices.

These instructions are either configured before (prepended) the system configures features defined in regular Firewall Management Center policies and settings, or after (appended). Any FlexConfig that depends on Secure Firewall Management Center-configured objects (for example, a network object) must be appended to the configuration deployment, or the needed objects would not be configured before the FlexConfig needed to refer to the objects.

For more information on configuring FlexConfig objects, see Configure FlexConfig Objects, on page 2690.

Geolocation

Each geolocation object you configure represents one or more countries or continents that the system has identified as the source or destination of traffic on your monitored network. You can use geolocation objects in various places in the system's web interface, including access control policies, SSL policies, remote access VPN, and event searches. For example, you could write an access control rule that blocks traffic to or from certain countries.

To ensure that you are using up-to-date information to filter your network traffic, Cisco strongly recommends that you regularly update your Geolocation Database (GeoDB).

Creating Geolocation Objects

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Geolocation** from the list of object types.

Step 3 Click **Add Geolocation**.

Step 4 Enter a **Name**.

Step 5 Check the check boxes for the countries and continents you want to include in your geolocation object. Checking a continent chooses all countries within that continent, as well as any countries that GeoDB updates may add under that continent in the future. Unchecking any country under a continent unchecks the continent. You can choose any combination of countries and continents.

Step 6 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Interface

Each interface can be assigned to a *security zone* and/or *interface group*. You then apply your security policy based on zones or groups. For example, you can assign the "inside" interface to the "inside" zone; and the "outside" interface to the "outside" zone. You can configure your access control policy to enable traffic to go from inside to outside, but not from outside to inside, for example. Some policies only support security zones, while other policies support zones and groups.

For more information about interface objects, see Security Zones and Interface Groups, on page 775.

To add interface objects, see Create Security Zone and Interface Group Objects, on page 777.

Key Chain

To enhance data security and protection of devices, rotating keys for authenticating IGP peers that have a duration of 180 days or less is introduced. The rotating keys prevent any malicious user from guessing the keys used for routing protocol authentication and thereby protecting the network from advertising incorrect routes and redirecting traffic. Changing the keys frequently reduces the risk of them eventually being guessed. When configuring authentication for routing protocols that provide key chains, configure the keys in a key chain to have overlapping lifetimes. This helps to prevent loss of key-secured communication due to absence of an active key. The rotating keys are applicable only for OSPFv2 protocol. If the key lifetime expires and no active keys are found, OSPF uses the last valid key to maintain the adjacency with peers.

Note Only MD5 cryptographic algorithm is used for authentication.

Lifetime of a Key

To maintain stable communications, each device stores key chain authentication keys and uses more than one key for a feature at the same time. Based on the send and accept lifetimes of a key, key chain management provides a secured mechanism to handle key rollover. The device uses the lifetimes of keys to determine which keys in a key chain are active.

Each key in a key chain has two lifetimes:

- Accept lifetime—The time interval within which the device accepts the key during key exchange with another device.
- Send lifetime—The time interval within which the device sends the key during key exchange with another device.

During a key send lifetime, the device sends routing update packets with the key. The device does not accept communication from other devices when the key sent is not within the accept lifetime of the key on the device.

If lifetimes are not configured then it is equivalent to configuring MD5 authentication key without timelines.

Key Selection

- When key chain has more than one valid key, OSPF selects the key that has the maximum life time.
- Key having an infinite lifetime is preferred.
- If keys have the same lifetime, then key with the higher key ID is preferred.

Creating Key Chain Objects

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Key Chain** from the list of object types.

Step 3 Click **Add Key Chain**.

Step 4 In the Add Key Chain Object dialog box, enter a name for the key chain in the **Name** field.

The name must start with an underscore or alphabet, followed by alphanumeric characters or special characters(-, _, +, .).

Step 5 To add a key to the key chain, click **Add**.

Step 6 Specify the key identifier in the **Key ID** field.

The key id value can be between 0 and 255. Use the value 0 only when you want to signal an invalid key.

Step 7 The **Algorithm** field and the **Crypto Encryption Type** field displays the supported algorithm and the encryption type, namely MD5 and Plain Text respectively.

Step 8 Enter the password in the **Crypto Key String** field, and re-enter the password in the **Confirm Crypto Key String** field.

- The password can be of a maximum length of 80 characters.
- The passwords cannot be a single digit nor those starting with a digit followed by a white space. For example, "0 pass" or "1" are invalid.

Step 9 To set the time interval for a device to accept/send the key during key exchange with another device, provide the lifetime values in the **Accept Lifetime** and **Send Lifetime** fields:

Note
The Date Time values default to UTC timezones.

The end time can be the duration, the absolute time when the accept/send lifetime ends, or never expires. The default end time is DateTime.

Following are the validation rules for the start and end values:

- Start lifetime cannot be null when the end lifetime is specified.
- The start lifetime for accept or send lifetime must be earlier than the respective end lifetime.

Step 10 Click **Add**.

Repeat steps 5 to 10 to create keys. Create a minimum of two keys for a key chain with overlapping lifetimes. This helps to prevent loss of key-secured communication due to absence of an active key.

Step 11 Manage overrides for the object:

- If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 1362.
- If you want to add override values to this object, expand the Override section and click **Add**; see Adding Object Overrides, on page 1362.

Step 12 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Network

A network object represents one or more IP addresses. You can use network objects and groups in various places, including access control policies, network variables, identity rules, network discovery rules, event searches, reports, identity policies, and so on.

When you configure an option that requires a network object, the list is automatically filtered to show only those objects that are valid for the option. For example, some options require host objects, while other options require subnets.

A network object can be one of the following types:

Host

A single IP address.

IPv4 example:

```
209.165.200.225
```

IPv6 example:

`2001:DB8::0DB8:800:200C:417A` or `2001:DB8:0:0:0DB8:800:200C:417A`

Range

A range of IP addresses.

IPv4 example:

```
209.165.200.225-209.165.200.250
```

IPv6 example:

```
2001:db8:0:cd30::1-2001:db8:0:cd30::1000
```

Network

An address block, also known as a subnet.

IPv4 example:

```
209.165.200.224/27
```

IPv6 example:

```
2001:DB8:0:CD30::/60
```

Note Security Intelligence ignores IP address blocks using a `/0` netmask.

FQDN

A single fully-qualified domain name (FQDN). You can limit FQDN resolution to IPv4 address only, IPv6 address only, or both IPv4 and IPv6 addresses. FQDNs must begin and end with a digit or letter. Only letters, digits, and hyphens are allowed as internal characters in an FQDN.

For example:

```
www.example.com
```

Note You can use FQDN objects in access control rules and prefilter rules, or manual NAT rules, only. The rules match the IP address obtained for the FQDN through a DNS lookup. To use an FQDN network object, ensure you have configured the DNS server settings in DNS Server Group, on page 1384 and the DNS platform settings in DNS, on page 949.

You *cannot* use FDQN network objects in identity rules.

Group

A group of network objects or other network object groups. You can create nested groups by adding one network object group to another network object group. You can nest up to 10 levels of groups.

Note You can add up to 100 network literals in a network object. Additionally, each nested network object group can contain a maximum of 100 network literals.

Network Wildcard Mask

You can create and manage wildcard mask objects from the Object Management page.

You can create network objects with expanded subnet IP address. The existing network object is extended to support both Network and Network Wildcard object. The network object using wildcard mask is listed as **Network Wildcard** against the **Type** column in the network object listing page.

A wildcard mask is an IP address that is a discontinuous mask of bits. You can use contiguous masks to create standard network objects and discontinuous masks for wildcard network objects.

Example IP Address	Network Wildcard?	Object Type
192.0.0.0/8	No	Network
10.10.0.0/255.255.0.0	No	Network
10.10.0.10/255.255.0.255	Yes	Network Wildcard
72.0.240.10/255.255.240.255	Yes	Network Wildcard

Note Network wildcard object and object group, which contains network wildcard objects, are allowed only while configuring the following policies:

- Prefilter policy
- Access control policy
- NAT policy

Guidelines and Limitations

- To create network wildcard objects, in the Firewall Management Center UI, choose **Objects** > **Object Management** > **Network** and click **Add Network** and then **Add Object**. Select the **Network** option and enter the value as expanded subnet mask. Example: 10.0.10.10/255.255.0.255.
- Object override, group object support, group object override, wildcard literals, and wildcard object import are supported.
- The network wildcard object is supported only for IPv4 addresses.
- The network wildcard object is supported from Firewall Management Center and Firewall Threat Defense 7.1 version onwards.
- Network wildcard objects are supported only for Snort-3.

Creating Network Objects

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Network** from the list of object types.

Step 3 Choose **Add Object** from the **Add Network** drop-down menu.

Step 4 Alternatively, you can clone an existing network object and edit the parameters to create a new network object. Click the **Clone** icon on the existing network object that you want to clone.

Step 5 Enter a **Name**.

Step 6 Optionally, enter a **Description**.

Step 7 In the **Network** field, select the required option and enter an appropriate value; see Network, on page 1401.

Note
You can add up to 100 network literals in a network object. Additionally, each nested network object group can contain a maximum of 100 network literals.

Step 8 (FQDN objects only) Select the DNS resolution from the **Lookup** drop-down menu to determine whether you want the IPv4, IPv6, or both IPv4 and IPv6 addresses associated with the FQDN.

Step 9 Manage overrides for the object:

- If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 1362.
- If you want to add override values to this object, expand the Override section and click **Add**; see Adding Object Overrides, on page 1362.

Step 10 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Importing Network Objects

For details on importing network objects, see Importing Objects, on page 1354.

PKI

PKI Objects for SSL Application

PKI objects represent the public key certificates and paired private keys required to support your deployment. Internal and trusted CA objects consist of certificate authority (CA) certificates; internal CA objects also contain the private key paired with the certificate. Internal and external certificate objects consist of server certificates; internal certificate objects also contain the private key paired with the certificate.

If you use trusted certificate authority objects and internal certificate objects to configure a connection to ISE/ISE-PIC, you can use ISE/ISE-PIC as an identity source.

If you use internal certificate objects to configure captive portal, the system can authenticate the identity of your captive portal device when connecting to users' web browsers.

If you use trusted certificate authority objects to configure realms, you can configure secure connections to LDAP or AD servers.

If you use PKI objects in SSL rules, you can match traffic encrypted with:

- the certificate in an external certificate object
- a certificate either signed by the CA in a trusted CA object, or within the CA's chain of trust

If you use PKI objects in SSL rules, you can decrypt:

- outgoing traffic by re-signing the server certificate with an internal CA object
- incoming traffic using the known private key in an internal certificate object

You can manually input certificate and key information, upload a file containing that information, or in some cases, generate a new CA certificate and private key.

When you view a list of PKI objects in the object manager, the system displays the certificate's Subject distinguished name as the object value. Hover your pointer over the value to view the full certificate Subject distinguished name. To view other certificate details, edit the PKI object.

Note The Firewall Management Center and managed devices encrypt all private keys stored in internal CA objects and internal certificate objects with a randomly generated key before saving them. If you upload private keys that are password protected, the appliance decrypts the key using the user-supplied password, then reencrypts it with the randomly generated key before saving it.

PKI Objects for Certificate Enrollment

A certificate enrollment object contains the Certification Authority (CA) server information and enrollment parameters that are required for creating Certificate Signing Requests (CSRs) and obtaining Identity Certificates from the specified CA. These activities occur in your Private Key Infrastructure (PKI).

The certificate enrollment object may also includes certificate revocation information. For more information on PKI, digital certificates, and certificate enrollment see PKI Infrastructure and Digital Certificates , on page 1505.

Internal Certificate Authority Objects

Each internal certificate authority (CA) object you configure represents the CA public key certificate of a CA your organization controls. The object consists of the object name, CA certificate, and paired private key. You can use internal CA objects and groups in SSL rules to decrypt outgoing encrypted traffic by re-signing the server certificate with the internal CA.

Note If you reference an internal CA object in a **Decrypt - Resign** SSL rule and the rule matches an encrypted session, the user's browser may warn that the certificate is not trusted while negotiating the SSL handshake. To avoid this, add the internal CA object certificate to either the client or domain list of trusted root certificates.

You can create an internal CA object in the following ways:

- import an existing RSA-based or elliptic curve-based CA certificate and private key
- generate a new self-signed RSA-based CA certificate and private key
- generate an unsigned RSA-based CA certificate and private key. You must submit a certificate signing request (CSR) to another CA to sign the certificate before using the internal CA object.

After you create an internal CA object containing a signed certificate, you can download the CA certificate and private key. The system encrypts downloaded certificates and private keys with a user-provided password.

Whether system-generated or user-created, you can modify the internal CA object name, but cannot modify other object properties.

You cannot delete an internal CA object that is in use. Additionally, after you edit an internal CA object used in an SSL policy, the associated access control policy goes out-of-date. You must re-deploy the access control policy for your changes to take effect.

CA Certificate and Private Key Import

You can configure an internal CA object by importing an X.509 v3 CA certificate and private key. You can upload files encoded in one of the following supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

If the private key file is password-protected, you can supply the decryption password. If the certificate and key are encoded in the PEM format, you can also copy and paste the information.

You can upload only files that contain proper certificate or key information, and that are paired with each other. The system validates the pair before saving the object.

Note If you configure a rule with the **Decrypt - Resign** action, the rule matches traffic based on the referenced internal CA certificate's encryption algorithm type, in addition to any configured rule conditions. You must upload an elliptic curve-based CA certificate to decrypt outgoing traffic encrypted with an elliptic curve-based algorithm, for example.

Importing a CA Certificate and Private Key

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Internal CAs**.

Step 3 Click **Import CA**.

Step 4 Enter a **Name**.

Step 5 Above the **Certificate Data** field, click **Browse** to upload a DER or PEM-encoded X.509 v3 CA certificate file.

Step 6 Above the **Key** field, click **Browse** to upload a DER or PEM-encoded paired private key file.

Step 7 If the uploaded file is password-protected, check the **Encrypted, and the password is:** check box, and enter the password.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Generating a New CA Certificate and Private Key

You can configure an internal CA object by providing identification information to generate a self-signed RSA-based CA certificate and private key.

The generated CA certificate is valid for ten years. The Valid From date is a week before generation.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Internal CAs**.

Step 3 Click **Generate CA**.

Step 4 Enter a **Name**.

Step 5 Enter the identification attributes.

Step 6 Click **Generate self-signed CA**.

New Signed Certificates

You can configure an internal CA object by obtaining a signed certificate from a CA. This involves two steps:

- Provide identification information to configure the internal CA object. This generates an unsigned certificate and paired private key, and creates a certificate signing request (CSR) to a CA you specify.
- After the CA issues the signed certificate, upload it to the internal CA object, replacing the unsigned certificate.

You can only reference an internal CA object in an SSL rule if it contains a signed certificate.

Creating an Unsigned CA Certificate and CSR

Procedure

Step 1 Choose **Objects** > **Object Management**.
Step 2 Expand the **PKI** node, and choose **Internal CAs**.
Step 3 Click **Generate CA**.
Step 4 Enter a **Name**.
Step 5 Enter the identification attributes.
Step 6 Click **Generate CSR**.
Step 7 Copy the CSR to submit to a CA.
Step 8 Click **OK**.

What to do next

- You must upload a signed certificate issued by a CA as described in Uploading a Signed Certificate Issued in Response to a CSR, on page 1407

Uploading a Signed Certificate Issued in Response to a CSR

Once uploaded, the signed certificate can be referenced in SSL rules.

Procedure

Step 1 Choose **Objects** > **Object Management**.
Step 2 Expand the **PKI** node, and choose **Internal CAs**.
Step 3 Click **Edit** (✎) next to the CA object containing the unsigned certificate awaiting the CSR.
Step 4 Click **Install Certificate**.

Step 5 Click **Browse** to upload a DER or PEM-encoded X.509 v3 CA certificate file.

Step 6 If the uploaded file is password protected, check the **Encrypted, and the password is:** check box, and enter the password.

Step 7 Click **Save** to upload a signed certificate to the CA object.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

CA Certificate and Private Key Downloads

You can back up or transfer a CA certificate and paired private key by downloading a file containing the certificate and key information from an internal CA object.

Caution Always store downloaded key information in a secure location.

The system encrypts the private key stored in an internal CA object with a randomly generated key before saving it to disk. If you download a certificate and private key from an internal CA object, the system first decrypts the information before creating a file containing the certificate and private key information. You must then provide a password the system uses to encrypt the downloaded file.

Caution Private keys downloaded as part of a system backup are decrypted, then stored in the unencrypted backup file.

Downloading a CA Certificate and Private Key

You can download CA certificates for both the current domain and ancestor domains.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Internal CAs**.

Step 3 Next to the internal CA object whose certificate and private key you want to download, click **Edit** ( ).

Step 4 Click **Download**.

Step 5 Enter an encryption password in the **Password** and **Confirm Password** fields.

Step 6 Click **OK**.

Trusted Certificate Authority Objects

Each trusted certificate authority (CA) object you configure represents a CA public key certificate belonging to a trusted CA. The object consists of the object name and CA public key certificate. You can use external CA objects and groups in:

- your SSL policy to control traffic encrypted with a certificate signed either by the trusted CA, or any CA within the chain of trust.
- your realm configurations to establish secure connections to LDAP or AD servers.
- your ISE/ISE-PIC connection. Select trusted certificate authority objects for the **pxGrid Server CA** and **MNT Server CA** fields.

After you create the trusted CA object, you can modify the name and add certificate revocation lists (CRL), but cannot modify other object properties. There is no limit on the number of CRLs you can add to an object. If you want to modify a CRL you have uploaded to an object, you must delete the object and recreate it.

Note Adding a CRL to an object has no effect when the object is used in your ISE/ISE-PIC integration configuration.

You cannot delete a trusted CA object that is in use. Additionally, after you edit a trusted CA object that is in use, the associated access control policy goes out-of-date. You must re-deploy the access control policy for your changes to take effect.

Trusted CA Object

You can configure an external CA object by uploading an X.509 v3 CA certificate. You can upload a file encoded in one of the following supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

If the file is password-protected, you must supply the decryption password. If the certificate is encoded in the PEM format, you can also copy and paste the information.

You can upload a CA certificate only if the file contains proper certificate information; the system validates the certificate before saving the object.

Adding a Trusted CA Object

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Trusted CAs**.

Step 3 Click **Add Trusted CAs**.

Step 4 Enter a **Name**.

Step 5 Click **Browse** to upload a DER or PEM-encoded X.509 v3 CA certificate file.

Step 6 If the file is password-protected, check the **Encrypted, and the password is:** check box, and enter the password.

Step 7 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Certificate Revocation Lists in Trusted CA Objects

You can upload CRLs to a trusted CA object. If you reference that trusted CA object in an SSL policy, you can control encrypted traffic based on whether the CA that issued the session encryption certificate subsequently revoked the certificate. You can upload files encoded in one of the following supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

After you add the CRL, you can view the list of revoked certificates. If you want to modify a CRL you have uploaded to an object, you must delete the object and recreate it.

You can upload only files that contain a proper CRL. There is no limit to the number of CRLs you can add to a trusted CA object. However, you must save the object each time you upload a CRL, before adding another CRL.

Note Adding a CRL to an object has no effect when the object is used in your ISE/ISE-PIC integration configuration.

Adding a Certificate Revocation List to a Trusted CA Object

Note Adding a CRL to an object has no effect when the object is used in your ISE/ISE-PIC integration configuration.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Trusted CAs**.

Step 3 Click **Edit** (✎) next to a trusted CA object.

If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Click **Add CRL** to upload a DER or PEM-encoded CRL file.

Step 5 Click **OK**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

External Certificate Objects

Each external certificate object you configure represents a server public key certificate that does not belong to your organization. The object consists of the object name and certificate. You can use external certificate objects and groups in SSL rules to control traffic encrypted with the server certificate. For example, you can upload a self-signed server certificate that you trust, but cannot verify with a trusted CA certificate.

You can configure an external certificate object by uploading an X.509 v3 server certificate. You can upload a file in one of the following supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

You can upload only files that contains proper server certificate information; the system validates the file before saving the object. If the certificate is encoded in the PEM format, you can also copy and paste the information.

Adding External Certificate Objects

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **External Certs**.

Step 3 Click **Add External Cert**.

Step 4 Enter a **Name**.

Step 5 Above the **Certificate Data** field, click **Browse** to upload a DER or PEM-encoded X.509 v3 server certificate file.

Step 6 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Internal Certificate Objects

Each internal certificate object you configure represents a server public key certificate belonging to your organization. The object consists of the object name, public key certificate, and paired private key. You can use internal certificate objects and groups in:

- your SSL rules to decrypt traffic incoming to one of your organization's servers using the known private key.
- your ISE/ISE-PIC connection. Select an internal certificate object for the **MC Server Certificate** field.
- your captive portal configuration to authenticate the identity of your captive portal device when connecting to users' web browsers. Select an internal certificate object for the **Server Certificate** field.

You can configure an internal certificate object by uploading an X.509 v3 RSA-based or elliptic curve-based server certificate and paired private key. You can upload a file in one of the following supported formats:

- Distinguished Encoding Rules (DER)
- Privacy-enhanced Electronic Mail (PEM)

If the file is password-protected, you must supply the decryption password. If the certificate and key are encoded in the PEM format, you can also copy and paste the information.

You can upload only files that contain proper certificate or key information, and that are paired with each other. The system validates the pair before saving the object.

After you create the internal certificate object, you can modify the name, but cannot modify other object properties.

You cannot delete an internal certificate object that is in use. Additionally, after you edit an internal certificate object that is in use, the associated access control policy goes out-of-date. You must re-deploy the access control policy for your changes to take effect.

Adding Internal Certificate Objects

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **PKI** node, and choose **Internal Certs**.

Step 3 Click **Add Internal Cert**.

Step 4 Enter a **Name**.

Step 5 Above the **Certificate Data** field, click **Browse** to upload a DER or PEM-encoded X.509 v3 server certificate file.

Step 6 Above the **Key** field, or click **Browse** to upload a DER or PEM-encoded paired private key file.

Step 7 If the uploaded private key file is password-protected, check the **Encrypted, and the password is:** check box, and enter the password.

Step 8 Click **Save**.

Certificate Enrollment Objects

Trustpoints let you manage and track CAs and certificates. A trustpoint is a representation of a CA or identity pair. A trustpoint includes the identity of the CA, CA-specific configuration parameters, and an association with one, enrolled identity certificate.

A certificate enrollment object contains the Certification Authority (CA) server information and enrollment parameters that are required for creating Certificate Signing Requests (CSRs) and obtaining Identity Certificates from the specified CA. These activities occur in your Private Key Infrastructure (PKI).

The certificate enrollment object may also includes certificate revocation information. For more information on PKI, digital certificates, and certificate enrollment see PKI Infrastructure and Digital Certificates , on page 1505.

How to Use Certificate Enrollment Objects

Certificate Enrollment Objects are used to enroll your managed devices into your PKI infrastructure, and create trustpoints (CA objects) on devices that support VPN connections by doing the following:

1. Define parameters for CA authentication and enrollment in a Certificate Enrollment Object. Specify shared parameters and use the override facility to specify unique object setting for different devices.
2. Associate and install this object on each managed device that requires the identity certificate. On the device, it becomes a *trustpoint*.

   When a certificate enrollment object is associated with and then installed on a device, the process of certificate enrollment starts immediately. The process is automatic for self-signed, SCEP, EST, and PKCS12 file enrollment types, meaning it does not require any additional administrator action. Manual certificate enrollment requires extra administrator action.
3. Specify the created trustpoint in your VPN configuration.

Managing Certificate Enrollment Objects

To manage certificate enrollment objects, go to **Objects** > **Object Management** > **PKI** > **Cert Enrollment** . The following information is shown:

- Existing certificate enrollment objects are listed in the **Name** column.

  Use the search field (the magnifying glass) to filter the list.
- The enrollment type of each object is shown in the **Type** column. The following enrollment methods can be used:
  - **Self Signed**—The managed device generates its own self signed root certificate.
  - **EST**—Enrollment over Secure Transport is used by the device to obtain an identity certificate from the CA.
  - **SCEP**—(Default) Simple Certificate Enrollment Protocol (SCEP) is used by the device to obtain an identity certificate from the CA.
  - **Manual**—The process of enrolling is carried out manually by the administrator.
  - **PKCS12 File**—Import a PKCS12 file on a Threat Defense managed device that supports VPN connectivity. A PKCS#12, or PFX or P12 file holds the server certificate, any intermediate certificates, and the private key in one encrypted file. Enter the **Passphrase** value for decryption.

- The **Override** column indicates whether the object allows overrides (a green check mark) or not (a red X). If a number is displayed, it is the number of overrides in place.

  Use the Override option to customize the object settings for each device that is part of the VPN configuration. Overriding makes each device's trustpoint details unique. Typically the Common Name or Subject is overridden for each device in the VPN configuration.

  See Object Overrides, on page 1361 for details and procedures on overriding objects of any type.

- **Edit** a previously created certificate enrollment object by clicking on the edit icon (a pencil). Editing can only be done if the enrollment object is not associated with any managed devices. Refer to the adding instructions for editing a certificate enrollment object. Failed enrollment objects can be edited.

- **Delete** a previously created certificate enrollment object by clicking on the delete icon (a trash can). You cannot delete a certificate enrollment object if it is associated with any managed device.

Press (+)**Add Cert Enrollment** to open the **Add Cert Enrollment** dialog and configure a Certificate Enrollment Object, see Adding Certificate Enrollment Objects, on page 1414. Then install the certificate on each managed, headend device.

Related Topics

Installing a Certificate Using Self-Signed Enrollment , on page 1489
Installing a Certificate using EST Enrollment, on page 1490
Installing a Certificate Using SCEP Enrollment, on page 1491
Installing a Certificate Using Manual Enrollment, on page 1491
Installing a Certificate Using a PKCS12 File, on page 1492

Adding Certificate Enrollment Objects

You can use these objects with Firewall Threat Defense devices. You must have Admin or Network Admin privileges to do this task.

Procedure

Step 1 Open the **Add Cert Enrollment** dialog:

- Directly from Object Management: In the **Objects** > **Object Management** > **PKI** > **Cert Enrollment** from the navigation pane, and press **Add Cert Enrollment**.
- While configuring a managed device: In the **Devices** > **Certificates** screen, choose Add **Add** > **Add New Certificate** and click (+) for the **Certificate Enrollment** field.

Step 2 Enter the **Name**, and optionally, a **Description** of this enrollment object.

When enrollment is complete, this name is the name of the trustpoint on the managed devices with which it is associated.

Step 3 Click the **CA Information** tab.

Step 4 Choose the **Enrollment Type**.

- **Self-Signed Certificate**—The managed device, acting as a CA, generates its own self-signed root certificate. No other information is needed in this pane.

Note
When enrolling a self-signed certificate you must specify the Common Name (CN) in the certificate parameters.

- **EST**—Enrollment over Secure Transport protocol. Specify the EST information. See Certificate Enrollment Object EST Options, on page 1416.
- **SCEP**—(Default) Simple Certificate Enrollment Protocol. Specify the SCEP information. See Certificate Enrollment Object SCEP Options, on page 1417.
- **Manual**
  - **CA Only**—Check this check box to create only the CA certificate from the selected CA. An identity certificate will not be created for this certificate.

    If you do not select this check box, a CA certificate is not mandatory. You can generate the CSR without having a CA certificate and obtain the identity certificate.
  - **CA Certificate**—Paste the CA certificate in the PEM format in the box. You can also obtain a CA certificate by copying it from another device.

    You can leave this box empty if you choose to generate a CSR without the CA certificate.
- **PKCS12 File**—Import a PKCS12 file on a Firewall Threat Defense managed device that supports VPN connectivity. A PKCS#12, or PFX, file holds a server certificate, intermediate certificates, and a private key in one encrypted file. Enter the **Passphrase** value for decryption.

Step 5 **Skip Check for CA flag in basic constraints of the CA Certificate**—Check this check box if you want to skip checking the basic constraints extension and the CA flag in a trustpoint certificate.

Step 6 **Validation Usage**—Choose from the options to validate the certificate during a VPN connection:

- **IPsec Client**—Validate an IPsec client certificate for a site-to-site VPN connection.
- **SSL Client**—Validate an SSL client certificate during a remote access VPN connection attempt.
- **SSL Server**—Select to validate an SSL server certificate, like as a Cisco Umbrella server certificate.

Step 7 (Optional) Click the **Certificate Parameters** tab and specify the certificate contents. See Certificate Enrollment Object Certificate Parameters, on page 1418.

This information is placed in the certificate and is readable by any party who receives the certificate from the router.

Step 8 (Optional) Click the **Key** tab and specify the Key information. See Certificate Enrollment Object Key Options, on page 1419.

Step 9 (Optional) Click the **Revocation** tab, and specify the revocation options: See Certificate Enrollment Object Revocation Options, on page 1420.

Step 10 **Allow Overrides** of this object if desired.

Whenever you modify the PKCS12 certificate enrollment object to permit overrides, you must update the **Passphrase** for the certificate on the device where it is overridden.

See Object Overrides, on page 1361 for a full description of object overrides.

Step 11 Click **Save**.

What to do next

Associate and install the enrollment object on a device to create a trustpoint on that device.

Related Topics

Installing a Certificate Using Self-Signed Enrollment , on page 1489
Installing a Certificate using EST Enrollment, on page 1490
Installing a Certificate Using SCEP Enrollment, on page 1491
Installing a Certificate Using Manual Enrollment, on page 1491
Installing a Certificate Using a PKCS12 File, on page 1492

Add Certificate Enrollment

Procedure

Step 1 Enter the **Name**.

Step 2 Paste the certificate information in the **IdP Certificate** field in PEM format.

Note
If the certificate is dependent on a root or intermediate certificate, you must install the dependant certificates. See Certificates, on page 1485.

Step 3 Click **Save**.

Certificate Enrollment Object EST Options

Secure Firewall Management Center Navigation Path

Objects > **Object Management** > **PKI** > **Cert Enrollment** . Click (+) **Add Cert Enrollment** to open the **Add Cert Enrollment** dialog, and select the **CA Information** tab.

Fields

Enrollment Type—set to **EST**.

Note

- EST enrollment type does not support EdDSA key.
- EST's ability to auto-enroll a device when its certificate expires is not supported.

Enrollment URL—The URL of the CA server to which devices should attempt to enroll.

Use an HTTPS URL in the form of **https://CA_name:port**, where *CA_name* is the host DNS name or IP address of the CA server. The *port* number is mandatory.

Username—The username to access the CA server.

Password / Confirm Password—The password to access the CA server.

Fingerprint—When retrieving the CA certificate using EST, you may enter the fingerprint for the CA server. Using the fingerprint to verify the authenticity of the CA server's certificate helps prevent an unauthorized party from substituting a fake certificate in place of the real one. Enter the **Fingerprint** for the CA server in hexadecimal format. If the value you enter does not match the fingerprint on the certificate, the certificate is rejected. Obtain the CA's fingerprint by contacting the server directly.

Source Interface—The interface that interacts with the CA server. By default, the diagnostic interface is displayed. To configure a data interface as the source interface, choose the respective security zone or interface group object.

Ignore EST Server Certificate Validations—The EST server certificate validation is done by default. Check the check box if you want to ignore Firewall Threat Defense validating EST server certificate.

Certificate Enrollment Object SCEP Options

Secure Firewall Management Center Navigation Path

Objects > **Object Management** > **PKI** > **Cert Enrollment** . Click (+) **Add Cert Enrollment** to open the **Add Cert Enrollment** dialog, and select the **CA Information** tab.

Fields

Enrollment Type—set to **SCEP**.

Enrollment URL—The URL of the CA server to which devices should attempt to enroll.

Use an HTTP URL in the form of `http://CA_name:port,` where CA_name is the host DNS name or IP address of the CA server. The port number is mandatory.

Note If the SCEP Server is referred with hostname/FQDN, configure DNS Server using FlexConfig object.

If the CA cgi-bin script location at the CA is not the default (/cgi-bin/pkiclient.exe), you must also include the nonstandard script location in the URL, in the form of http://CA_name:port/script_location, where script_location is the full path to the CA scripts.

Challenge Password / Confirm Password—The password used by the CA server to validate the identity of the device. You can obtain the password by contacting the CA server directly or by entering the following address in a web browser: `http://URLHostName/certsrv/mscep/mscep.dll`. The password is good for 60 minutes from the time you obtain it from the CA server. Therefore, it is important that you deploy the password as soon as possible after you create it.

Retry Period—The interval between certificate request attempts, in minutes. Value can be 1 to 60 minutes. The default is 1 minute.

Retry Count—The number of retries that should be made if no certificate is issued upon the first request. Value can be 1 to 100. The default is 10.

CA Certificate Source—Specify how the CA certificate will be obtained.

- **Retrieve Using SCEP** (Default, and only supported option)—Retrieve the certificate from the CA server using the Simple Certificate Enrollment Process (SCEP). Using SCEP requires a connection between your device and the CA server. Ensure there is a route from your device to the CA server before beginning the enrollment process.

Fingerprint—When retrieving the CA certificate using SCEP, you may enter the fingerprint for the CA server. Using the fingerprint to verify the authenticity of the CA server's certificate helps prevent an unauthorized party from substituting a fake certificate in place of the real one. Enter the **Fingerprint** for the CA server in hexadecimal format. If the value you enter does not match the fingerprint on the certificate, the certificate is rejected. Obtain the CA's fingerprint by contacting the server directly, or by entering the following address in a web browser: **`http://<URLHostName>/certsrv/mscep/mscep.dll`**.

Certificate Enrollment Object Certificate Parameters

Specify additional information in certificate requests sent to the CA server. This information is placed in the certificate and can be viewed by any party who receives the certificate from the router.

Secure Firewall Management Center Navigation Path

Objects > **Object Management** > **PKI** > **Cert Enrollment** . Press (+) **Add Cert Enrollment** to open the **Add Cert Enrollment** dialog, and select the **Certificate Parameters** tab.

Fields

Enter all information using the standard LDAP X.500 format.

- **Include FQDN**—Whether to include the device's fully qualified domain name (FQDN) in the certificate request. Choices are:
  - **Use Device Hostname as FQDN**
  - **Don't use FQDN in certificate**
  - **Custom FQDN**—Select this and then specify it in the **Custom FQDN** field that displays.
- **Include Device's IP Address**—The interface whose IP address is included in the certificate request.
- **Common Name (CN)**—The X.500 common name to include in the certificate.

Note When enrolling a self-signed certificate you must specify the Common Name (CN) in the certificate parameters.

- **Organization Unit (OU)**—The name of the organization unit (for example, a department name) to include in the certificate.
- **Organization (O)**—The organization or company name to include in the certificate.
- **Locality (L)**—The locality to include in the certificate.
- **State (ST)**—The state or province to include in the certificate.
- **County Code (C)**—The country to include in the certificate. These codes conform to ISO 3166 country abbreviations, for example "US" for the United States of America.
- **Email (E)**—The email address to include in the certificate.
- **Include Device's Serial Number**—Whether to include the serial number of the device in the certificate. The CA uses the serial number to either authenticate certificates or to later associate a certificate with a particular device. If you are in doubt, include the serial number, as it is useful for debugging purposes.

Certificate Enrollment Object Key Options

Secure Firewall Management Center Navigation Path

Objects > **Object Management** > **PKI** > **Cert Enrollment** . Press (+) **Add Cert Enrollment** to open the **Add Cert Enrollment** dialog, and select the **Key** tab.

Fields

- **Key Type**—RSA, ECDSA, EdDSA.

Note

- For EST enrollment type, do not select EdDSA key as it is not supported.
- EdDSA is supported only in Site-to-Site VPN topologies.
- EdDSA is not supported as an identity certificate for the Remote Access VPN.

- **Key Name**—If the key pair you want to associate with the certificate already exists, this field specifies the name of that key pair. If the key pair does not exist, this field specifies the name to assign to the key pair that will be generated during enrollment. If you do not specify a name, the fully qualified domain name (FQDN) key pair is used instead.
- **Key Size**—If the key pair does not exist, defines the desired key size (modulus), in bits. The recommended size is 2048 bits. The larger the modulus size, the more secure the key. However, keys with larger modulus sizes take longer to generate (a minute or more when larger than 512 bits) and longer to process when exchanged.

Important

- On Firewall Management Center and Firewall Threat Defense Versions 7.0 and higher, you cannot enroll certificates with RSA key sizes smaller than 2048 bits and keys using SHA-1 with the RSA Encryption algorithm. However, you can use PKI Enrollment of Certificates with Weak-Crypto to allow certificates that use SHA-1 with RSA Encryption algorithm and smaller key size.
- You cannot generate RSA keys with sizes smaller than 2048 bits for Firewall Threat Defense 7.0, even when you enable the weak-crypto option.

- **Advanced Settings**—Select **Ignore IPsec Key Usage** if you do not want to validate values in the key usage and extended key usage extensions of IPsec remote client certificates. You can suppress key usage checking on IPsec client certificates. By default this option is not enabled.

Note For site-to-site VPN connection, if you use a Windows Certificate Authority (CA), the default Application Policies extension is **IP security IKE intermediate**. If you are using this default setting, you must select the **Ignore IPsec Key Usage** option for the object you select. Otherwise, the endpoints cannot complete the site-to-site VPN connection.

PKI Enrollment of Certificates with Weak-Crypto

SHA-1 hashing signature algorithm, and RSA key sizes that are smaller than 2048 bits for certification are not supported on Firewall Management Center and Firewall Threat Defense Version 7.0 and higher. You cannot enroll certificates with RSA key sizes that are smaller than 2048 bits.

To override these restrictions on Firewall Management Center 7.0 managing Firewall Threat Defenses running Versions lesser than 7.0, you can use the enable weak-crypto option on the Firewall Threat Defense. We do not recommend you to permit weak-crypto keys, because, such keys are not as secure as the ones with higher key sizes.

Note Firewall Threat Defense 7.0 or higher does not support generating RSA keys with sizes smaller than 2048 bits even when you permit weak-crypto.

To enable weak-crypto on the device, navigate to the **Devices** > **Certificates** page. Click the **Enable Weak-Crypto**( ) button provided against the Firewall Threat Defense device. When the weak-crypto option is enabled, the button changes to . By default, the weak-crypto option is disabled.

Note When a certificate enrollment fails due to weak cipher usage, the Firewall Management Center displays a warning message prompting you to enable the weak-crypto option. Similarly, when you turn on the enable weak-crypto button, the Firewall Management Center displays a warning message before enabling weak-crypto configuration on the device.

Upgrading Earlier Versions to Firewall Threat Defense 7.0

When you are upgrading to Firewall Threat Defense 7.0, the existing certificate configurations are retained. However, if those certificates have RSA keys smaller than 2048 bits and use SHA-1 encryption algorithm, they cannot be used to establish VPN connections. You must either procure a certificate with RSA key sizes bigger than 2048 bits or enable the permit weak-crypto option for VPN connections.

Certificate Enrollment Object Revocation Options

Specify whether to check the revocation status of a certificate by choosing and configuring the method. Revocation checking is off by default, neither method (CRL or OCSP) is checked.

Secure Firewall Management Center Navigation Path

Objects > **Object Management** > **PKI** > **Cert Enrollment** . Press ( ) **Add Cert Enrollment** to open the **Add Cert Enrollment** dialog, and select the **Revocation** tab.

Fields

- **Enable Certificate Revocation Lists**—Check to enable CRL checking.
  - **Use CRL distribution point from the certificate**—Check to obtain the revocation lists distribution URL from the certificate.
  - **Use static URL configured**—Check this to add a static, pre-defined distribution URL for revocation lists. Then add the URLs.

    CRL Server URLs—The URL of the LDAP server from which the CRL can be downloaded.

    URLs must start with **`http://`**. Include a port number in the URL. Enclose IPv6 addresses in square brackets, for example: *http://[0:0:0:0:0.18:0a01:7c16]*.
- **Enable Online Certificate Status Protocol (OCSP)**—Check to enable OCSP checking.

  OCSP Server URL—The URL of the OCSP server checking for revocation if you require OCSP checks.

  URLs must start with **`http://`**. Enclose IPv6 addresses in square brackets, for example: *http://[0:0:0:0:0.18:0a01:7c16]*.
- **Consider the certificate valid if revocation information cannot be reached**—Checked by default. Uncheck if you do not want to allow this.

Policy List

Use the Configure Policy List page to create, copy, and edit policy list policy objects. You can create policy list objects to use when you are configuring route maps. When a policy list is referenced within a route map, all of the match statements within the policy list are evaluated and processed. Two or more policy lists can be configured with a route map. A policy list can also coexist with any other preexisting match and set statements that are configured within the same route map but outside of the policy list. When multiple policy lists perform matching within a route map entry, all policy lists match on the incoming attribute only.

You can use this object with Firewall Threat Defense devices.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **Policy List** from the table of contents.

Step 2 Click **Add Policy List**.

Step 3 Enter a name for the policy list object in the **Name** field. Object names are not case-sensitive.

Step 4 Select whether to allow or block access for matching conditions from the **Action** drop-down list.

Step 5 Click the **Interface** tab to distribute routes that have their next hop out of one of the interfaces specified.

In the **Zones/Interfaces** list, add the zones that contain the interfaces through which the device communicates with the management station. For interfaces not in a zone, you can type the interface name into the field below the **Selected Zone/Interface** list and click **Add**. The host will be configured on a device only if the device includes the selected interfaces or zones.

Step 6 Click the **Address** tab to redistribute any routes that have a destination address that is permitted by a standard access list or prefix list.

Choose whether to use an **Access List** or **Prefix List** for matching and then enter or select the Standard Access List Objects or Prefix list objects you want to use for matching.

Step 7 Click the **Next Hop** tab to redistribute any routes that have a next hop router address passed by one of the access lists or prefix lists specified.

Choose whether to use an **Access List** or **Prefix List** for matching and then enter or select the Standard Access List Objects or Prefix list objects you want to use for matching.

Step 8 Click the **Route Source** tab to redistribute routes that have been advertised by routers and access servers at the address specified by the access lists or prefix list.

Choose whether to use an **Access List** or **Prefix List** for matching and then enter or select the Standard Access List Objects or Prefix list objects you want to use for matching.

Step 9 Click the **AS Path** tab to match a BGP autonomous system path. If you specify more than one AS path, then the route can match either AS path.

Step 10 Click the **Community Rule** tab to enable matching of the BGP community or extended community with the specified community list objects or the extended community list objects respectively. If you specify more than one rule, the routes are verified against the rules until a matching permit or deny is met.

a) To specify a community list to the rule, click **Edit** (✎) given in the **Selected Community List** field. The community lists appear under **Available Community List**. Select the required list, click **Add**, and then click **Ok**.

To enable matching the BGP community exactly with the specified community, check the **Match the specified community exactly** check box.

b) To add the extended community list, click **Edit** (✎) given in the **Selected Extended Community List** field. The extended community lists appear under the **Available Extended Community List**. Select the required list, click **Add**, and then click **Ok**.

Note
The extended community lists are applicable only for configuring import or export of routes.

Step 11 Click the **Metric & tag** tab to match the metric and security group tag of a route.

a) Enter the metric values to use for matching in the **Metric** field. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified metric. The metric values can range from 0 to 4294967295.

b) Enter the tag values to use for matching in the **Tag** field. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified security group tag. The tag values can range from 0 to 4294967295.

Step 12 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 1362.

Step 13 Click **Save**.

Port

Port objects represent different protocols in slightly different ways:

TCP and UDP

A port object represents the transport layer protocol, with the protocol number in parentheses, plus an optional associated port or port range. For example: `TCP(6)/22`.

ICMP and ICMPv6 (IPv6-ICMP)

A port object represents the Internet layer protocol plus an optional type and code. For example: `ICMP(1):3:3`.

You can restrict an ICMP or IPV6-ICMP port object by type and, if applicable, code. For more information on ICMP types and codes, see:

- http://www.iana.org/assignments/icmp-parameters/icmp-parameters.xml
- http://www.iana.org/assignments/icmpv6-parameters/icmpv6-parameters.xml

Other

A port object can represent other protocols that do not use ports.

The system provides default port objects for well-known ports. You cannot modify or delete these default objects. You can create custom port objects in addition to the default objects.

You can use port objects and groups in various places in the system's web interface, including access control policies, identity rules, network discovery rules, port variables, and event searches. For example, if your organization uses a custom client that uses a specific range of ports and causes the system to generate excessive and misleading events, you can configure your network discovery policy to exclude monitoring those ports.

When using port objects, observe the following guidelines:

- You cannot add any protocol other than TCP or UDP for source port conditions in access control rules. Also, you cannot mix transport protocols when setting both source and destination port conditions in a rule.
- If you add an unsupported protocol to a port object group used in a source port condition, the rule where it is used does not take affect on the managed device when the configuration is deployed.
- If you create a port object containing both TCP and UDP ports, then add it as a source port condition in a rule, you cannot add a destination port, and vice versa.

Creating Port Objects

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Port** from the list of object types.

Step 3 Choose **Add Object** from the **Add Port** drop-down list.

Alternatively, you can clone an existing port object and edit the parameters to create a new port object. Click the **Clone** icon on the existing port object that you want to clone.

Step 4 Enter a **Name**.

Step 5 Choose a **Protocol**.

Step 6 Depending on the protocol you chose, constrain by **Port**, or choose an ICMP **Type** and **Code**.

You can enter ports from `1` to `65535`. Use a hyphen to specify a port range. You must constrain the object by port if you chose to match **All** protocols, using the **Other** drop-down list.

Step 7 Manage overrides for the object:

- If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 1362.
- If you want to add override values to this object, expand the Override section and click **Add**; see Adding Object Overrides, on page 1362.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Importing Port Objects

For details on importing port objects, see Importing Objects, on page 1354.

Prefix List

You can create prefix list objects for IPv4 and IPv6 to use when you are configuring route maps, policy maps, OSPF Filtering, or BGP Neighbor Filtering.

Configure IPv6 Prefix List

Use the Configure IPv6 Prefix list page to create, copy and edit prefix list objects. You can create prefix list objects to use when you are configuring route maps, policy maps, OSPF Filtering, or BGP Neighbor Filtering.

You can use this object with Firewall Threat Defense devices.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **Prefix Lists** > **IPv6 Prefix List** from the table of contents.

Step 2 Click **Add Prefix List**.

Step 3 Enter a name for the prefix list object in the **Name** field on the **New Prefix List Object** window.

Step 4 Click **Add** on the **New Prefix List Object** window.

Step 5 Select the appropriate action, Allow or Block from the **Action** drop-down list, to indicate the redistribution access.

Step 6 Enter a unique number that indicates the position a new prefix list entry will have in the list of prefix list entries already configured for this object, in the **Sequence No.** field. If left blank, the sequence number will default to five more than the largest sequence number currently in use.

Step 7 Specify the IPv6 address in the IP address/mask length format in the **IP address** field. The mask length must be a valid value between 1-128.

Step 8 Enter the minimum prefix length in the **Minimum Prefix Length** field. The value must be greater than the mask length and less than or equal to the Maximum Prefix Length, if specified.

Step 9 Enter the maximum prefix length in the **Maximum Prefix Length** field. The value must be greater than or equal to the Minimum Prefix Length, if present, or greater than the mask length if the Minimum Prefix Length is not specified.

Step 10 Click **Add**.

Step 11 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 1362.

Step 12 Click **Save**.

Configure IPv4 Prefix List

Use the Configure IPv4 Prefix list page to create, copy and edit prefix list objects. You can create prefix list objects to use when you are configuring route maps, policy maps, OSPF Filtering, or BGP Neighbor Filtering.

You can use this object with Firewall Threat Defense devices.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **Prefix Lists** > **IPv4 Prefix List** from the table of contents.

Step 2 Click **Add Prefix List**.

Step 3 Enter a name for the prefix list object in the **Name** field on the **New Prefix List Object** window.

Step 4 Click **Add**.

Step 5 Select the appropriate action, Allow or Block from the **Action** drop-down list, to indicate the redistribution access.

Step 6 Enter a unique number that indicates the position a new prefix list entry will have in the list of prefix list entries already configured for this object, in the **Sequence No.** field. If left blank, the sequence number will default to five more than the largest sequence number currently in use.

Step 7 Specify the IPv4 address in the IP address/mask length format in the **IP address** field. The mask length must be a valid value between 1- 32.

Step 8 Enter the minimum prefix length in the **Minimum Prefix Length** field. The value must be greater than the mask length and less than or equal to the Maximum Prefix Length, if specified.

Step 9 Enter the maximum prefix length in the **Maximum Prefix Length** field. The value must be greater than or equal to the Minimum Prefix Length, if present, or greater than the mask length if the Minimum Prefix Length is not specified.

Step 10 Click **Add**.

Step 11 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 1362.

Step 12 Click **Save**.

Route Map

Route maps are used when redistributing routes into any routing process. They are also used when generating a default route into routing process. A route map defines which of the routes from the specified routing protocol are allowed to be redistributed into the target routing process. Configure a route map, to create a new route map entry for a Route Map object or to edit an existing one.

You can use this object with Firewall Threat Defense devices.

Before you begin

A Route Map may use one or mores of these objects; it is not mandatory to add all these objects. Create and use any of these objects as required, to configure your route map.

- Add ACLs.
- Add Prefix Lists.
- Add AS Path.
- Add Community Lists.
- Add Extended Community Lists.

Note The extended community lists are applicable only for configuring import or export of routes.

- Add Policy Lists.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **Route Map** from the table of contents.

Step 2 Click **Add Route Map**.

Step 3 Click **Add** on the **New Route Map Object** window.

Step 4 In the **Sequence No.** field, enter a number, from 0 through 65535, that indicates the position a new route map entry has in the list of route maps entries already configured for this route map object.

Note
We recommend that you number clauses in intervals of at least 10 to reserve numbering space in case you want to insert clauses in the future.

Step 5 Select the appropriate action, Allow or Block, from the **Redistribution** drop-down list, to indicate the redistribution access.

Step 6 Click the **Match Clauses** tab to match (routes/traffic) based on the following criteria, which you select in the table of contents:

- **Security Zones** —Match traffic based on the (ingress/egress) interfaces. You can select zones and add them, or type in interface names and add them.

- **IPv4** — Match IPv4 (routes/traffic) based on the following criteria; select the tab to define the criteria.
  a. Click the **Address** tab to match routes based on the route address. For IPv4 addresses, choose whether to use an Access list or Prefix list for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching.
  b. Click the **Next Hop** tab to match routes based on the next hop address of a route. For IPv4 addresses, choose whether to use an access list or Prefix list for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching.
  c. Click the **Route Source** tab to match routes based on the advertising source address of the route. For IPv4 addresses, choose whether to use an access list or Prefix list for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching.
- **IPv6** —Match IPv6 (routes/traffic) based on the route address, next-hop address or advertising source address of route.
- **BGP** —Match BGP (routes/traffic) based on the following criteria; select the tab to define the criteria.
  a. Click the **AS Path** tab to enable matching the BGP autonomous system path access list with the specified path access list. If you specify more than one path access list, then the route can match either path access list.
  b. Click the **Community List** tab to enable matching of the BGP community or extended community with the specified community list objects or the extended community list objects respectively.
     - To specify a community list to the rule, click **Edit** (✎) given in the **Selected Community List** field. The community lists appears under **Available Community List**. Select the required list, click **Add**, and then click **Ok**. For information on how to create community list objects, see Community List, on page 1379
     - To add the extended community list, click **Edit** (✎) given in the **Selected Extended Community List** field. The extended community lists appears under the **Available Extended Community List**. Select the required list, click **Add**, and then click **Ok**. For information on how to create extended community list objects, see Extended Community, on page 1380.

     To enable matching the BGP community exactly with the specified community list objects, check the **Match the specified community exactly** check box. This option is not applicable for the extended community list.

     Note
     If you specify more than one rule, the routes are verified against the rules until a matching permit or deny condition is met. Any route that does not match at least one Match community will not be advertised for outbound route maps.

  c. Click the **Policy List** tab to configure a route map to evaluate and process a BGP policy. When multiple policy lists perform matching within a route map entry, all policy lists match on the incoming attribute only.
- **Others** —Match routes or traffic based on the following criteria.
  a. Enter the metric values to use for matching in the **Metric Route Value** field, to enable matching the metric of a route. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified metric. The metric values can range from 0 to 4294967295.

b. Enter the tag values to use for matching in the **Tag Values** field. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified security group tag. The tag values can range from 0 to 4294967295.

c. Check the appropriate **Route Type** option to enable matching of the route type. Valid route types are External1, External2, Internal, Local, NSSA-External1, and NSSA-External2. You can choose more than one route type from the list.

Step 7 Click the **Set Clauses** tab to set routes/traffic based on the following criteria, which you select in the table of contents:

- **Metric Values**—Set either Bandwidth, all of the values or none of the values.

  a. Enter a metric value or bandwidth in Kbits per second in the **Bandwidth** field. Valid values are an integer value in the range from 0 to 4294967295.

  b. Select to specify the type of metric for the destination routing protocol, from the **Metric Type** drop-down list. Valid values are : internal, type-1, or type-2.

- **BGP Clauses** —Set BGP routes based on the following criteria; select the tab to define the criteria.

  a. Click the **AS Path** tab to modify an autonomous system path for BGP routes.

  1. Enter an AS path number in the **Prepend AS Path** field to prepend an arbitrary autonomous system path string to BGP routes. Usually the local AS number is prepended multiple times, increasing the autonomous system path length. If you specify more than one AS path number then the route can prepend either AS number.
  2. Enter an AS path number in the **Prepend Last AS to AS Path** field to prepend the AS path with the last AS number. Enter a value for the AS number from 1 to 10.
  3. Check the **Convert route tag into AS path** check box to convert the tag of a route into an autonomous system path.

  b. Click the **Community List** tab to set the community attributes:

  Under **Specific Community**:

  1. Click the **None** radio button, to remove the community attribute from the prefixes that pass the route map.
  2. Click the **Specific Community** radio button, to enter a community number, if applicable. Valid values are from 1 to 4294967295.
  3. Check the **Add to existing communities** check box, to add the community to the already existing communities.
  4. Select the **Internet**, **No-Advertise**, or **No-Export** check-boxes to use one of the well-known communities.

  Under **Specific Extended Community**, in the **Route Target** field, enter the route target number in *ASN:nn* format:

  - You can enter values that ranges from 1:1 to 65534:65535.

    You can add a single route target or a set of route targets separated by commas in a single entry. For example, *1:2,1:4,1:6*.

- You can have a maximum of 8 route targets in an entry.
- You cannot have redundant route target entries across route maps.

c. Click the **Others** tab to set additional attributes.

1. Check the **Set Automatic Tag** check-box to automatically compute the tag value.
2. Enter a preference value for the autonomous system path in the **Set Local Preference** field. Enter a value between 0 and 4294967295.
3. Enter a BGP weight for the routing table in the **Set Weight** field. Enter a value between 0 and 65535.
4. Select to specify the BGP origin code. Valid values are **Local IGP** Local IGP and **Incomplete**.
5. In the IPv4 Settings section, specify a next hop IPv4 address of the next hop to which packets are output. It need not be an adjacent router. If you specify more than one IPv4 address then the packets can output at either IP address.

   Select to specify an IPv4 prefix list in the **Prefix List** drop-down list.
6. In the IPv6 Settings section, specify a next hop IPv6 address of the next hop to which packets are output. It need not be an adjacent router. If you specify more than one IPv6 address, then the packets can output at any of the IP addresses.

   Select to specify an IPv6 prefix in the **Prefix List** drop-down list.

Step 8 Click **Add**.

Step 9 If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 1362.

Step 10 Click **Save**.

Security Intelligence

Security Intelligence functionality requires the IPS license (for Firewall Threat Defense devices) or the Protection license (all other device types).

Security Intelligence *lists* and *feeds* are collections of IP addresses, domain names, and URLs that you can use to quickly filter traffic that matches an entry on a list or feed.

- A list is a static collection that you manage manually.
- A feed is a dynamic collection that updates on an interval over HTTP or HTTPS.

Security Intelligence lists/feeds are grouped into:

- DNS (Domain names)
- Network (IP addresses)
- URLs

System-Provided Feeds

Cisco provides the following feeds as Security Intelligence objects:

- Security Intelligence feeds updated regularly with the latest threat intelligence from Talos:
  - Cisco-DNS-and-URL-Intelligence-Feed (under DNS Lists and Feeds)
  - Cisco-Intelligence-Feed (for IP addresses, under Network Lists and Feeds)

  You cannot delete the system-provided feeds, but you can change the frequency of (or disable) their updates.
- Cisco-TID-Feed (under Network Lists and Feeds)

  This feed is not used in the Security Intelligence tab of the access control policy.

  Instead, you must enable and configure Secure Firewall Threat Intelligence Director to use this feed, which is a collection of TID observables data.

  Use this object to set how frequently this data is published to TID elements.

  For more information, see Secure Firewall Threat Intelligence Director, on page 2735.

Predefined Lists: Global Block Lists and Global Do Not Block Lists

The system ships with predefined global Block lists and Do Not Block lists for domains (DNS), IP addresses (Networks), and URLs.

These lists are empty until you populate them. To build these lists, see Global and Domain Security Intelligence Lists, on page 1431.

By default, access control and DNS policies use these lists as part of Security Intelligence.

Custom Feeds

You can use third-party feeds, or use a custom internal feed to easily maintain an enterprise-wide Block list in a large deployment with multiple Secure Firewall Management Center appliances.

See Custom Security Intelligence Feeds, on page 1437.

Custom Lists

Custom lists can augment and fine-tune feeds and the Global lists.

See Custom Security Intelligence Lists, on page 1438.

Where Security Intelligence Lists and Feeds Are Used

- IP address and address blocks—Use Block and Do Not Block lists in access control policies, as part of Security Intelligence.
- Domain Names—Use Block and Do Not Block lists in DNS policies, as part of Security Intelligence.
- URLs—Use Block and Do Not Block lists in access control policies, as part of Security Intelligence. You can also use URL lists in access control and QoS rules, whose analysis and traffic handling phases occur after Security Intelligence.

How to Modify Security Intelligence Objects

To add or delete entries on a Block list, Do Not Block list, feed, or sinkhole object:

Object Type	Edit Capabilities	Requires Redeploy After Edit?
Custom Block and Do Not Block lists	Upload new and replacement lists using the object manager.	No
Default (but custom-populated) Block lists and Do Not Block lists: Global, descendant, and domain-specific	Add entries using the context menu or delete entries using the object manager.	No
System-provided Intelligence Feeds	Disable or change update frequency using the object manager.	No
Custom feeds	Fully modify using the object manager.	No
Sinkhole	Fully modify using the object manager.	Yes

Global and Domain Security Intelligence Lists

The Firewall Management Center includes Global Block and Do-Not-Block lists, which enable you to use Security Intelligence to consistently block specific connections or exempt certain connections from being blocked, allowing them to be evaluated by other threat detection processes you have configured.

For example, if you notice a set of routable IP addresses in intrusion events associated with exploit attempts, you can immediately block those IP addresses. Although it may take a few minutes for your changes to propagate, you do not have to redeploy.

By default, Access control and DNS policies use these Global lists, which apply to all security zones. You can opt not to use these lists on a per-policy basis.

Note These options apply to Security Intelligence only. Security Intelligence cannot block traffic that has already been fastpathed. Similarly, adding an item to a Security Intelligence Do Not Block list does not automatically trust or fastpath matching traffic. For more information, see About Security Intelligence, on page 1885.

Security Intelligence Lists and Multitenancy

Multitenancy adds:

- Domain lists—Block or Do Not Block lists whose contents apply to a particular subdomain only. The Global lists are Domain lists for the Global domain.
- Descendant Domain lists—Block or Do Not Block lists that aggregate the Domain lists of the current domain's descendants.

Domain Lists

In addition to being able to access (but not edit) the Global lists, each subdomain has its own named lists, the contents of which apply only to that subdomain. For example, a subdomain named Company A owns:

- Domain Block list - Company A and Domain Do Not Block list - Company A
- Domain Block list for DNS - Company A, Domain Do Not Block list for DNS - Company A
- Domain Block list for URL - Company A, Domain Do Not Block list for URL - Company A

Any administrator at or above the current domain can populate these lists. You can use the context menu to add an item to the Block or Do Not Block list in the current and all descendant domains. However, only an administrator in the associated domain can remove an item from a Domain list.

For example, a Global administrator could choose to add the same IP address to the Block list in the Global domain and Company A's domain, but not add it to the Block list in Company B's domain. This action would add the same IP address to:

- Global Block list (where it can be removed only by Global administrators)
- Domain Block list - Company A (where it can be removed only by Company A administrators)

Descendant Domain Lists

A Descendant Domain list is a Do Not Block list or Block list that aggregates the Domain lists of the current domain's descendants. Leaf domains do not have Descendant Domain lists.

Descendant Domain lists are useful because a higher-level domain administrator can enforce general Security Intelligence settings, while still allowing subdomain users to add items to a Block or Do Not Block list in their own deployment.

For example, the Global domain has the following Descendant Domain lists:

- Descendant Block lists - Global, Descendant Do Not Block lists - Global
- Descendant Block lists for DNS - Global, Descendant Do Not Block lists for DNS - Global
- Descendant Block lists for URL - Global, Descendant Do Not Block lists for URL - Global

Note Descendant Domain lists do not appear in the object manager because they are symbolic aggregations, not hand-populated lists. They appear where you can use them: in access control and DNS policies.

Add Entries to Global Security Intelligence Lists

When reviewing events and dashboards, you can instantly block future traffic involving IP addresses, domains, and URLs that appear in those events by adding them to a predefined Block list.

Similarly, if Security Intelligence is blocking traffic that you want evaluated by threat detection processes subsequent to Security Intelligence blocking, you can add IP addresses, domains, and URLs from events to a predefined Do Not Block list.

Traffic is evaluated against entries on these lists during the Security Intelligence phase of threat detection.

For more information about these lists, see Global and Domain Security Intelligence Lists, on page 1431.

Before you begin

Because adding an entry to a Security Intelligence list affects access control, you must have one of the following user roles:

- Administrator
- A combination of roles: Network Admin or Access Admin, plus Security Analyst and Security Approver
- A custom role with both Modify Access Control Policy and Deploy Configuration to Devices permissions

If appropriate, verify that these lists are used in the policies in which you expect them to be used.

Procedure

Step 1 Navigate to an event that includes an IP address, domain, or URL that you want to always block using Security Intelligence, or exempt from Security Intelligence blocking.

Step 2 Right-click the IP address, domain, or URL and choose the appropriate option:

Item Type	Context Menu Option
IP address	Add IP to Block List Add IP to Do-Not-Block List These options add the IP address to the respective lists for Networks.
URL	Add URL to Global Block List for URL Add URL to Global Do-Not-Block List for URL
Domain of a URL in the URL field	Add Domain to Global Block List for URL Add Domain to Global Do-Not-Block List for URL
Domain in the DNS Query field	Add Domain to Global Block List for DNS Add Domain to Global Do-Not-Block List for DNS

What to do next

You do NOT need to redeploy for these changes to take effect.

If you want to delete an item from a list, see Delete Entries from Global Security Intelligence Lists, on page 1433.

Delete Entries from Global Security Intelligence Lists

Note To add entries to these lists, see Add Entries to Global Security Intelligence Lists, on page 1432.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Click **Security Intelligence**.

Step 3 Click the appropriate option:

- **Network Lists and Feeds** (for IP addresses)
- **DNS Lists and Feeds** (for domain names)
- **URL Lists and Feeds**

Step 4 Click the pencil beside the Global Block or Global Do-Not-Block list.

Step 5 Click the trash button beside the entry to delete.

List and Feed Updates for Security Intelligence

List and feed updates replace the existing list or feed file with the contents of the new file. Contents of existing and new files are not merged.

If the system downloads a corrupt feed or a feed with no recognizable entries, the system continues using the old feed data (unless it is the first download). However, if the system can recognize even one entry in the feed, it uses the entries it can recognize.

By default, each feed updates the Management Center every two hours; you can modify this frequency. Any updates the Management Center receives are passed immediately to managed devices. In addition, managed devices poll the Firewall Management Center every 30 minutes for changes. You cannot modify this frequency.

To modify feed update intervals, see Changing the Update Frequency for Security Intelligence Feeds, on page 1434.

Changing the Update Frequency for Security Intelligence Feeds

You can specify the intervals at which the Firewall Management Center updates Security Intelligence Feeds.

For details about feed updates, see List and Feed Updates for Security Intelligence, on page 1434.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **Security Intelligence** node, then choose the feed type whose frequency you want to change.

The system-provided URL feed is combined with the domain feed under **DNS Lists and Feeds**.

Step 3 Next to the feed you want to update, click **Edit** ( ).

If **View** ( ) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 4 Edit the **Update Frequency**.

Step 5 Click **Save**.

Custom Security Intelligence Lists and Feeds

Custom Lists and Feeds: Requirements

List and Feed Formatting

Each list or feed must be a simple text file no larger than 500MB. List files must have the .txt extension. Include one entry or comment per line: one IP address, one URL, one domain name.

Tip The number of entries you can include is limited by the maximum size of the file. For example, a URL list with no comments and an average URL length of 100 characters (including Punycode or percent Unicode representations and newlines) can contain more than 5.24 million entries.

In a DNS list entry, you can specify an asterisk (`*`) wildcard character for a domain label. All labels match the wildcard. For example, an entry of `www.example.*` matches both `www.example.com` and `www.example.co`.

If you add comment lines within the source file, they must start with the pound (`#`) character. If you upload a source file with comments, the system removes your comments during upload. Source files you download contain all your entries without your comments.

Feed Requirements

When you configure a feed, you specify its location using a URL; the URL cannot be Punycode-encoded.

For feed update intervals of 30 minutes or less, you must specify an MD5 URL. This prevents frequent downloads of unchanged feeds. If your feed server does not provide an MD5 URL, you must use a download interval of at least 30 minutes.

If you use an MD5 checksum, the checksum must be stored in a simple text file with only the checksum. Comments are not supported.

URL Lists and Feeds: URL Syntax and Matching Criteria

Security Intelligence URL lists and feeds, including custom lists and feeds and entries in the global Block list and Do Not Block list, can include the following, which have the matching behavior as described:

- Hostnames

  For example, **`www.example.com`**.

- URLs

  `example.com` matches **`example.com`** and all subdomains, including **`www.example.com`**, **`eu.example.com`**, **`example.com/abc`**, and **`www.example.com/def`** -- but NOT **`example.co.uk`** or **`examplexyz.com`** or **`example.com.malicious-site.com`**

  When a URL feed or list includes a single entry, every URL that ends with those domains is identified and blocked.

Example: When the following URL feed —**`www.netflix.com, www.amazon.*, org, edu, www.hulu.*`** is added to the global Block List, the following contents are blocked:http://www.amazon.in, http://www.rajiv.org/, http://www.edu.edu/, http://www.edu.org/, http://org.org/ and http://edu.edu.

You can also include an entire URL path, such as **`https://www.cisco.com/c/en/us/products/security/firewalls/index.html`**

Note You can create a custom URL, Network, and DNS feeds, wherein, you can add the username and password inside the URL itself, for example: **`https://admin:password@server.domain.com/list.txt`**

However, if your password contain special characters such as a colon (:) or an at sign (@), the transmission would fail. Ensure that your password does not have any special characters. Alternatively, you could use an encoded password in the URL.

- A slash at the end of a URL to specify an exact match

  `example.com/` matches ONLY **`example.com`**; it does NOT match **`www.example.com`** or any other URL.

- A wildcard (*) to represent any domain in a URL

  An asterisk can represent a complete domain string separated by dots, but not a partial domain string, and not any part of the URL following the first slash.

  Valid examples:

  - **`*.example.com`**
  - **`www.*.com`**
  - **`example.*`**

    (This will match **`example.com`** and **`example.org`** and **`example.de`**, for example, but NOT **`example.co.uk`**)
  - **`*.example.*`**
  - **`example.*/`**

  Invalid examples:

  - **`example*.com`**
  - **`example.com/*`**

- IP addresses (IPv4)

  For IPv6 addresses, or to use ranges or CIDR notation, use the Security Intelligence Network object.

  You can include one or more wildcards representing an octet, for example 10.10.10.* or 10.10.*.*.

See also Custom Security Intelligence Lists, on page 1438.

Custom Security Intelligence Feeds

Custom or third-party Security Intelligence feeds allow you to augment the system-provided Intelligence Feeds with other regularly-updated reputable Block lists and Do Not Block lists on the Internet. You can also set up an internal feed, which is useful if you want to update multiple Secure Firewall Management Center appliances in your deployment using one source list.

Note You cannot add address blocks to Block or Do Not Block lists using a `/0` netmask in a Security Intelligence feed. If you want to monitor or block all traffic targeted by a policy, use an access control rule with the **Monitor** or **Block** rule action, respectively, and a default value of `any` for the **Source Networks** and **Destination Networks**.

You also can configure the system to use an MD5 checksum to determine whether to download an updated feed. If the checksum has not changed since the last time the system downloaded the feed, the system does not need to re-download it. You may want to use MD5 checksums for internal feeds, especially if they are large.

Note The system does **not** perform peer SSL certificate verification when downloading custom feeds, nor does the system support the use of certificate bundles or self-signed certificates to verify the remote peer.

If you want strict control over when the system updates a feed from the Internet, you can disable automatic updates for that feed. However, automatic updates ensure the most up-to-date, relevant data.

Manually updating Security Intelligence feeds updates all feeds, including the Intelligence Feeds.

See complete requirements at Custom Lists and Feeds: Requirements, on page 1435.

Creating Security Intelligence Feeds

You must have the IPS license (for Firewall Threat Defense devices) or the Protection license (all other device types).

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **Security Intelligence** node, then choose a feed type you want to add.

Step 3 Click the option appropriate to the feed type you chose above:

- **Add Network Lists and Feeds** (for IP addresses)
- **Add DNS Lists and Feeds**
- **Add URL Lists and Feeds**

Step 4 Enter a **Name** for the feed.

Step 5 Choose **Feed** from the **Type** drop-down list.

Step 6 Enter a **Feed URL**.

Step 7 Enter an **MD5 URL**.

This is used to determine whether the feed contents have changed since the last update, so the system does not download unchanged feeds.

MD5 URL is required for update intervals shorter than 30 minutes.

If your feed server does not provide an MD5 URL, you must choose an interval of at least 30 minutes.

Step 8 Choose an **Update Frequency**.

Step 9 Click **Save**.
Unless you disabled feed updates, the system attempts to download and verify the feed.

Manually Updating Security Intelligence Feeds

You must have the IPS license (for Firewall Threat Defense devices) or the Protection license (all other device types).

Before you begin

At least one device must already be added to the management center.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **Security Intelligence** node, then choose a feed type.

Step 3 Click **Update Feeds**, then confirm.

Step 4 Click **OK**.

After the Secure Firewall Management Center downloads and verifies the feed updates, it communicates any changes to its managed devices. Your deployment begins filtering traffic using the updated feeds.

Custom Security Intelligence Lists

Security Intelligence lists are simple static lists of IP addresses and address blocks, URLs, or domain names that you manually upload to the system. Custom lists are useful if you want to augment and fine-tune feeds or one of the global lists, for a single Secure Firewall Management Center's managed devices.

For example, if a reputable feed improperly blocks your access to vital resources but is overall useful to your organization, you can create a custom Do Not Block list that contains only the improperly classified IP addresses, rather than removing the IP address feed object from the access control policy's Block list.

Note You cannot add address blocks to a Block or Do Not Block list using a `/0` netmask in a Security Intelligence list. If you want to monitor or block all traffic targeted by a policy, use an access control rule with the **Monitor** or **Block** rule action, respectively, and a default value of `any` for the **Source Networks** and **Destination Networks**.

Regarding list entry formatting, note the following:

- Netmasks for address blocks can be integers from `0` to `32` or `0` to `128`, for IPv4 and IPv6, respectively.
- Unicode in domain names must be encoded in Punycode format, and are case insensitive.
- Characters in domain names are case-insensitive.
- Unicode in URLs should be encoded in percent-encoding format.
- Characters in URL subdirectories are case-sensitive.
- List entries that start with the pound sign (`#`) are treated as comments.
- See additional formatting requirements at Custom Lists and Feeds: Requirements, on page 1435.

Regarding matching list entries, note the following:

- The system matches sub-level domains if a higher-level domain exists in a URL or DNS list. For example, if you add `example.com` to a DNS list, the system matches both `www.example.com` and `test.example.com`.
- The system does not perform DNS lookups (forward or reverse) on DNS or URL list entries. For example, if you add `http://192.168.0.2` to a URL list, and it resolves to `http://www.example.com`, the system only matches `http://192.168.0.2`, not `http://www.example.com`.

Uploading New Security Intelligence Lists to the Secure Firewall Management Center

To modify a Security Intelligence list, you must make your changes to the source file and upload a new copy. You cannot modify the file's contents using the web interface. If you do not have access to the source file, download a copy from the system.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **Security Intelligence** node, then choose a list type.

Step 3 Click the option appropriate to the list you chose above:

- **Add Network Lists and Feeds** (for IP addresses)
- **Add DNS Lists and Feeds**
- **Add URL Lists and Feeds**

Step 4 Enter a **Name**.

Step 5 From the **Type** drop-down list, choose **List**.

Step 6 Click **Browse** to browse to the list `.txt` file, then click **Upload**.

Step 7 Click **Save**.

What to do next

You do not need to redeploy these changes to take effect. If you want to delete an entry from the list, see Delete Entries from Global Security Intelligence Lists, on page 1433.

Updating Security Intelligence Lists

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Expand the **Security Intelligence** node, then choose a list type.

Step 3 Next to the list you want to update, click **Edit** ( ).

If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 If you need a copy of the list to edit, click **Download**, then follow your browser's prompts to save the list as a text file.

Step 5 Make changes to the list as necessary.

Step 6 On the Security Intelligence pop-up window, click **Browse** to browse to the modified list, then click **Upload**.

Step 7 Click **Save**.

What to do next

You do not need to redeploy these changes to take effect. If you want to delete an entry from the list, see Delete Entries from Global Security Intelligence Lists, on page 1433.

Sinkhole

A sinkhole object represents either a DNS server that gives non-routable addresses for all domain names within the sinkhole, or an IP address that does not resolve to a server. You can reference the sinkhole object within a DNS policy rule to redirect matching traffic to the sinkhole. You must assign the object both an IPv4 address and an IPv6 address.

Creating Sinkhole Objects

You must have the IPS license (for Firewall Threat Defense devices) or the Protection license (all other device types).

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Sinkhole** from the list of object types.

Step 3 Click **Add Sinkhole**.

Step 4 Enter a **Name**.

Step 5 Enter the **IPv4 Address** and **IPv6 Address** of your sinkhole.

Step 6 You have the following options:

- If you want to redirect traffic to a sinkhole server, choose **Log Connections to Sinkhole**.
- If you want to redirect traffic to a non-resolving IP address, choose **Block and Log Connections to Sinkhole**.

Step 7 If you want to assign an Indication of Compromise (IoC) type to your sinkhole, choose one from the **Type** drop-down.

Step 8 Click **Save**.

SLA Monitor

Each Internet Protocol Service Level Agreement (SLA) monitor defines a connectivity policy to a monitored address and tracks the availability of a route to the address. The route is periodically checked for availability by sending ICMP echo requests and waiting for the response. If the requests time out, the route is removed from the routing table and replaced with a backup route. SLA monitoring jobs start immediately after deployment and continue to run unless you remove the SLA monitor from the device configuration (that is, they do not age out). The Internet Protocol Service Level Agreement (SLA) Monitor Object is used in the Route Tracking field of an IPv4 Static Route Policy. IPv6 routes do not have the option to use SLA monitor via route tracking.

You can use these objects with Firewall Threat Defense devices.

Procedure

Step 1 Select **Objects** > **Object Management** and choose **SLA Monitor** from the table of contents.

Step 2 Click **Add SLA Monitor**.

Step 3 Enter a name for the object in the **Name** field.

Step 4 (Optional) Enter a description for the object in the **Description** field.

Step 5 Enter the frequency of ICMP echo request transmissions, in seconds, in the **Frequency** field. Valid values range from 1 to 604800 seconds (7 days). The default is 60 seconds.

Note
The frequency cannot be less than the timeout value; you must convert frequency to milliseconds to compare the values.

Step 6 Enter the ID number of the SLA operation in the **SLA Monitor ID** field. Values range from 1 to 2147483647. You can create a maximum of 2000 SLA operations on a device. Each ID number must be unique to the policy and the device configuration.

Step 7 Enter the amount of time that must pass after an ICMP echo request before a rising threshold is declared, in milliseconds, in the **Threshold** field. Valid values range from 0 to 2147483647 milliseconds. The default is 5000 milliseconds. The threshold value is used only to indicate events that exceed the defined value. You can use these events to evaluate the proper timeout value. It is not a direct indicator of the reachability of the monitored address.

Note
The threshold value should not exceed the timeout value.

Step 8 Enter the amount of time that the SLA operation waits for a response to the ICMP echo requests, in milliseconds, in the **Timeout** field. Values range from 0 to 604800000 milliseconds (7 days). The default is 5000 milliseconds. If a response is not received from the monitored address within the amount of time defined in this field, the static route is removed from the routing table and replaced by the backup route.

Note
The timeout value cannot exceed the frequency value (adjust the frequency value to milliseconds to compare the numbers).

Step 9 Enter the size of the ICMP request packet payload, in bytes, in the **Data Size** field. Values range from 0 to 16384 bytes. The default is 28 bytes, which creates a total ICMP packet of 64 bytes. Do not set this value higher than the maximum allowed by the protocol or the Path Maximum Transmission Unit (PMTU). For purposes of reachability, you might need to increase the default data size to detect PMTU changes between the source and the target. A low PMTU can affect session performance and, if detected, might indicate that the secondary path should be used.

Step 10 Enter a value for type of service (ToS) defined in the IP header of the ICMP request packet in the **ToS** field. Values range from 0 to 255. The default is 0. This field contains information such as delay, precedence, reliability, and so on. It can be used by other devices on the network for policy routing and features such as committed access rate.

Step 11 Enter the number of packets that are sent in the **Number of Packets** field. Values range from 1 to 100. The default is 1 packet.

Note
Increase the default number of packets if you are concerned that packet loss might falsely cause the Secure Firewall Threat Defense device to believe that the monitored address cannot be reached.

Step 12 Enter the IP address that is being monitored for availability by the SLA operation, in the **Monitored Address** field.

Step 13 The **Available Zones** list displays both zones and interface groups. In the **Zones/Interfaces** list, add the zones or interface groups that contain the interfaces through which the device communicates with the management station. To specify a single interface, you need to create a zone or the interface groups for the interface; see Create Security Zone and Interface Group Objects, on page 777. The host will be configured on a device only if the device includes the selected interfaces or zones.

Step 14 Click **Save**.

Time Range

Use time range objects to define time periods that you will use to determine when rules apply.

Note Time-based ACLs is supported in Snort 3 also from Firewall Management Center 7.0 onwards.

Creating Time Range Objects

If you want a policy to apply only during a specified time range, create a time range object, then specify that object in the policy. Note that this object works on Firewall Threat Defense devices only.

You can specify time range objects only in policy types listed at the bottom of this topic.

Note The timezone represents the device's local time and is used ONLY for applying the time ranges in rules in the policies that support the time ranges. The timezone does not change the configured time of the device. To verify the configuration, in the Firewall Threat Defense CLI, use the **show time-range timezone** and **show time** commands (see the Cisco Secure Firewall Threat Defense Command Reference guide). In addition, the timezone of a chassis overrides the management center timezone.

Before you begin

Time ranges are applied based on the time zone associated with the device that processes the traffic. By default, this is UTC. To change the time zone associated with a device, go to **Devices** > **Platform Settings**.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Time Range** from the list of object types.

Step 3 Click **Add Time Range**.

Step 4 Enter values.

Observe the following guidelines:

- If you see a red error box around the object name you have entered, mouse over the **Name** field to see naming restrictions.
- All times are in UTC, unless you specify a time zone for the device in **Device** > **Platform Settings**.
- Enter times using a 24-hour clock. For example, enter 1:30 PM as 13:30.
- To specify a single continuous range, such as typical weekend hours (Fridays at 5pm through Mondays at 8am, including evenings and nights), choose Range Type **Range**.
- To specify part of multiple days, such as Monday through Friday from 8am to 5pm (excluding evenings, nights, and early mornings every day), choose Range Type **Daily Interval**.
- You can specify up to 28 time periods in a single object.
- To specify multiple noncontiguous times of day or different hours for different days, create multiple recurring intervals. For example, to apply a policy at all times other than standard working hours, create a single time range object with the following two recurring intervals:
  - A Daily Interval for Monday through Friday from 5pm through 8am, and
  - A Range recurring interval for Friday at 5pm through Monday at 8am.

Step 5 Click **Save**.

What to do next

Configure time ranges in any of the following:

- Access control rules
- Prefilter rules
- Tunnel rules
- VPN group policy

In a VPN group policy object, specify the time range object using the **Access Hours** field. For details, see Configure Group Policy Objects, on page 1466 and Group Policy Advanced Options, on page 1472.

Time Zone

To specify a local time zone for a managed device, create a time zone object and specify it in the device platform settings policy assigned to the device.

This device local time is used ONLY for applying time ranges in rules in policies that support time ranges, such as access control, prefilter, and VPN Group policies. If you do not assign a time zone to a device, UTC is used by default when applying time ranges in these policies. No other functionality in the system uses the time zone specified in a time zone object.

Time zone objects are supported only for Firewall Threat Defense devices.

Note Time-based ACLs is supported in Snort 3 also from Firewall Management Center 7.0 onwards.

Tunnel Zone

A *tunnel zone* represents certain types of plaintext, passthrough tunnels that you explicitly tag for special analysis. A tunnel zone is not an interface object, even though you can use it as an interface constraint in some configurations.

For detailed information, see Tunnel Zones and Prefiltering, on page 1928.

URL

Important For best practices for using this and similar options in Security Intelligence configurations and for URL rules in access control and QoS policies, see Manual URL Filtering Options, on page 1872.

A URL object defines a single URL or IP address, whereas a URL group object can define more than one URL or address. You can use URL objects and groups in various places in the system's web interface, including access control policies and event searches.

When creating URL objects, keep the following points in mind:

- If you do not include a path (that is, there are no / characters in the URL), the match is based on the server's hostname only. If you include one or more / character, the entire URL string is used for a substring match. Then, a URL is considered a match if any of the following are true:
  - The string is at the beginning of the URL.
  - The string follows a dot.
  - The string contains a dot in the beginning.
  - The string follows the :// characters.

  For example, ign.com matches ign.com or www.ign.com, but not versign.com.

Note We recommend that you do not use manual URL filtering to block or allow individual web pages or parts of sites (that is, URL strings with / characters), as servers can be reorganized and pages moved to new paths.

- The system disregards the encryption protocol (HTTP vs HTTPS). In other words, if you block a website, both HTTP and HTTPS traffic to that website is blocked, unless you use an application condition to target a specific protocol. When creating a URL object, you do not need to specify the protocol when creating an object. For example, use example.com rather than http://example.com.
- If you plan to use a URL object to match HTTPS traffic in an access control rule, create the object using the subject common name in the public key certificate used to encrypt the traffic. Also, the system disregards subdomains within the subject common name, so do not include subdomain information. For example, use example.com rather than www.example.com.

  However, please understand that the subject common name in the certificate might be completely unrelated to a web site's domain name. For example, the subject common name in the certificate for youtube.com is *.google.com (this of course might change at any time). You will get more consistent results if you use the SSL Decryption policy to decrypt HTTPS traffic so that URL filtering rules work on decrypted traffic.

Note URL objects will not match HTTPS traffic if the browser resumes a TLS session because the certificate information is no longer available. Thus, even if you carefully configure the URL object, you might get inconsistent results for HTTPS connections.

Creating URL Objects

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **URL** from the list of object types.

Step 3 Choose **Add Object** from the **Add URL** drop-down list.

Step 4 Enter a **Name**.

Step 5 Optionally, enter a **Description**.

Step 6 Enter the **URL** or IP address.

Step 7 Manage overrides for the object:

- If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 1362.
- If you want to add override values to this object, expand the Override section and click **Add**; see Adding Object Overrides, on page 1362.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Variable Set

Variables represent values commonly used in intrusion rules to identify source and destination IP addresses and ports. You can also use variables in intrusion policies to represent IP addresses in rule suppressions, adaptive profile updates, and dynamic rule states.

Tip Preprocessor rules can trigger events regardless of the hosts defined by network variables used in intrusion rules.

You use variable sets to manage, customize, and group your variables. You can use the default variable set provided by the system or create your own custom sets. Within any set you can modify predefined default variables and add and modify user-defined variables.

Most of the shared object rules and standard text rules that the system provides use predefined default variables to define networks and port numbers. For example, the majority of the rules use the variable `$HOME_NET` to specify the protected network and the variable `$EXTERNAL_NET` to specify the unprotected (or outside) network. In addition, specialized rules often use other predefined variables. For example, rules that detect exploits against web servers use the `$HTTP_SERVERS` and `$HTTP_PORTS` variables.

Rules are more effective when variables more accurately reflect your network environment. At a minimum, you should modify default variables in the default set. By ensuring that a variable such as `$HOME_NET` correctly defines your network and `$HTTP_SERVERS` includes all web servers on your network, processing is optimized and all relevant systems are monitored for suspicious activity.

To use your variables, you link variable sets to intrusion policies associated with access control rules or with the default action of an access control policy. By default, the default variable set is linked to all intrusion policies used by access control policies.

Adding a variable to any set adds it to all sets; that is, each variable set is a collection of all variables currently configured on your system. Within any variable set, you can add user-defined variables and customize the value of any variable.

Initially, the system provides a single, default variable set comprised of predefined default values. Each variable in the default set is initially set to its default value, which for a predefined variable is the value set by the Talos Intelligence Group and provided in rule updates.

Although you can leave predefined default variables configured to their default values, Cisco recommends that you modify a subset of predefined variables.

You could work with variables only in the default set, but in many cases you can benefit most by adding one or more custom sets, configuring different variable values in different sets, and perhaps even adding new variables.

When using multiple sets, it is important to remember that the *current value* of any variable in the default set determines the *default value* of the variable in all other sets.

When you select **Variable Sets** on the Object Manager page, the object manager lists the default variable set and any custom sets you created.

On a freshly installed system, the default variable set is comprised only of the default variables predefined by Cisco.

Each variable set includes the default variables provided by the system and all custom variables you have added from any variable set. Note that you can edit the default set, but you cannot rename or delete the default set.

Caution Importing an access control or an intrusion policy overwrites existing default variables in the default variable set with the imported default variables. If your existing default variable set contains a custom variable not present in the imported default variable set, the unique variable is preserved.

Related Topics

Managing Variables, on page 1458
Managing Variable Sets, on page 1457

Variable Sets in Intrusion Policies

By default, the system links the default variable set to all intrusion policies used in an access control policy. When you deploy an access control policy that uses an intrusion policy, intrusion rules that you have enabled in the intrusion policy use the variable values in the linked variable set.

When you modify a custom variable set used by an intrusion policy in an access control policy, the system reflects the status for that policy as out-of-date on the Access Control Policy page. You must re-deploy the access control policy to implement changes in your variable set. When you modify the default set, the system reflects the status of all access control policies that use intrusion policies as out-of-date, and you must re-deploy all access control policies to implement your changes.

Variables

Variables belong to one of the following categories:

Default Variables

Variables provided by the system. You cannot rename or delete a default variable, and you cannot change its default value. However, you can create a customized version of a default variable.

Customized Variables

Variables you create. These variables can include:

- *customized default variables*

  When you edit the value for a default variable, the system moves the variable from the Default Variables area to the Customized Variables area. Because variable values in the default set determine the default values of variables in custom sets, customizing a default variable in the default set modifies the default value of the variable in all other sets.

- *user-defined variables*

  You can add and delete your own variables, customize their values within different variable sets, and reset customized variables to their default values. When you reset a user-defined variable, it remains in the Customized Variables area.

  User-defined variables can be one of the following types:

  - *network* variables specify the IP addresses of hosts in your network traffic.
  - *port* variables specify TCP or UDP ports in network traffic, including the value `any` for either type.

For example, if you create custom standard text rules, you might also want to add your own user-defined variables to more accurately reflect your traffic or as shortcuts to simplify the rule creation process. Alternatively, if you create a rule that you want to inspect traffic in the "demilitarized zone" (or DMZ) only, you can create a variable named `$DMZ` whose value lists the server IP addresses that are exposed. You can then use the `$DMZ` variable in any rule written for this zone.

Advanced Variables

Variables provided by the system under specific conditions. These variables have a very limited deployment.

Predefined Default Variables

By default, the system provides a single default variable set, which is comprised of predefined default variables. The Talos Intelligence Group uses rule updates to provide new and updated intrusion rules and other intrusion policy elements, including default variables.

Because many intrusion rules provided by the system use predefined default variables, you should set appropriate values for these variables. Depending on how you use variable sets to identify traffic on your network, you can modify the values for these default variables in any or all variable sets.

Caution Importing an access control or an intrusion policy overwrites existing default variables in the default variable set with the imported default variables. If your existing default variable set contains a custom variable not present in the imported default variable set, the unique variable is preserved.

The following table describes the variables provided by the system and indicates which variables you typically would modify. For assistance determining how to tailor variables to your network, contact Professional Services or Support.

Table 76: System-Provided Variables

Variable Name	Description	Modify?
`$AIM_SERVERS`	Defines known AOL Instant Messenger (AIM) servers, and is used in chat-based rules and rules that look for AIM exploits.	Not required.
`$DNS_SERVERS`	Defines Domain Name Service (DNS) servers. If you create a rule that affects DNS servers specifically, you can use the `$DNS_SERVERS` variable as a destination or source IP address.	Not required in current rule set.
`$EXTERNAL_NET`	Defines the network that the system views as the unprotected network, and is used in many rules to define the external network.	Yes, you should adequately define `$HOME_NET` and then exclude `$HOME_NET` as the value for `$EXTERNAL_NET`.
`$FILE_DATA_PORTS`	Defines non-encrypted ports used in intrusion rules that detect files in a network stream.	Not required.
`$FTP_PORTS`	Defines the ports of FTP servers on your network, and is used for FTP server exploit rules.	Yes, if your FTP servers use ports other than the default ports (you can view the default ports in the web interface).
`$GTP_PORTS`	Defines the data channel ports where the packet decoder extracts the payload inside a GTP (General Packet Radio Service [GPRS] Tunneling Protocol) PDU.	Not required.
`$HOME_NET`	Defines the network that the associated intrusion policy monitors, and is used in many rules to define the internal network.	Yes, to include the IP addresses for your internal network.
`$HTTP_PORTS`	Defines the ports of web servers on your network, and is used for web server exploit rules.	Yes, if your web servers use ports other than the default ports (you can view the default ports in the web interface).
`$HTTP_SERVERS`	Defines the web servers on your network. Used in web server exploit rules.	Yes, if you run HTTP servers.
`$ORACLE_PORTS`	Defines Oracle database server ports on your network, and is used in rules that scan for attacks on Oracle databases.	Yes, if you run Oracle servers.
`$SHELLCODE_PORTS`	Defines the ports you want the system to scan for shell code exploits, and is used in rules that detect exploits that use shell code.	Not required.
`$SIP_PORTS`	Defines the ports of SIP servers on your network, and is used for SIP exploit rules.	Not required.

Variable Name	Description	Modify?
$SIP_SERVERS	Defines SIP servers on your network, and is used in rules that address SIP-targeted exploits.	Yes, if you run SIP servers, you should adequately define $HOME_NET and then include $HOME_NET as the value for $SIP_SERVERS.
$SMTP_SERVERS	Defines SMTP servers on your network, and is used in rules that address exploits that target mail servers.	Yes, if you run SMTP servers.
$SNMP_SERVERS	Defines SNMP servers on your network, and is used in rules that scan for attacks on SNMP servers.	Yes, if you run SNMP servers.
$SNORT_BPF	Identifies a legacy advanced variable that appears only when it existed on your system in a software release before Version 5.3.0 that you subsequently upgraded to Version 5.3.0 or greater.	No, you can only view or delete this variable. You cannot edit it or recover it after deleting it.
$SQL_SERVERS	Defines database servers on your network, and is used in rules that address database-targeted exploits.	Yes, if you run SQL servers.
$SSH_PORTS	Defines the ports of SSH servers on your network, and is used for SSH server exploit rules.	Yes, if your SSH servers use ports other than the default port (you can view the default ports in the web interface).
$SSH_SERVERS	Defines SSH servers on your network, and is used in rules that address SSH-targeted exploits.	Yes, if you run SSH servers, you should adequately define $HOME_NET and then include $HOME_NET as the value for $SSH_SERVERS.
$TELNET_SERVERS	Defines known Telnet servers on your network, and is used in rules that address Telnet server-targeted exploits.	Yes, if you run Telnet servers.
$USER_CONF	Provides a general tool that allows you to configure one or more features not otherwise available via the web interface. Conflicting or duplicate $USER_CONF configurations will halt the system.	No, only as instructed in a feature description or with the guidance of Support.

Network Variables

Network variables represent IP addresses you can use in intrusion rules that you enable in an intrusion policy and in intrusion policy rule suppressions, dynamic rule states, and adaptive profile updates. Network variables differ from network objects and network object groups in that network variables are specific to intrusion policies and intrusion rules, whereas you can use network objects and groups to represent IP addresses in various places in the system's web interface, including access control policies, network variables, intrusion rules, network discovery rules, event searches, reports, and so on.

You can use network variables in the following configurations to specify the IP addresses of hosts on your network:

- intrusion rules—Intrusion rule **Source IPs** and **Destination IPs** header fields allow you to restrict packet inspection to the packets originating from or destined to specific IP addresses.

- suppressions—The **Network** field in source or destination intrusion rule suppressions allows you to suppress intrusion event notifications when a specific IP address or range of IP addresses triggers an intrusion rule or preprocessor.
- dynamic rule states—The **Network** field in source or destination dynamic rule states allows you to detect when too many matches for an intrusion rule or preprocessor rule occur in a given time period.
- adaptive profile updates—When you enable adaptive profile updates, the adaptive profiles **Networks** field identifies hosts where you want to improve reassembly of packet fragments and TCP streams in passive deployments.

When you use variables in the fields identified in this section, the variable set you link to an intrusion policy determines the variable values in the network traffic handled by an access control policy that uses the intrusion policy.

You can add any combination of the following network configurations to a variable:

- any combination of network variables, network objects, and network object groups that you select from the list of available networks
- individual network objects that you add from the New Variable or Edit Variable page, and can then add to your variable and to other existing and future variables
- literal, single IP addresses or address blocks

  You can list multiple literal IP addresses and address blocks by adding each individually. You can list IPv4 and IPv6 addresses and address blocks alone or in any combination. When specifying IPv6 addresses, you can use any addressing convention defined in RFC 4291.

The default value for included networks in any variable you add is the word `any`, which indicates any IPv4 or IPv6 address. The default value for excluded networks is `none`, which indicates no network. You can also specify the address `::` in a literal value to indicate any IPv6 address in the list of included networks, or no IPv6 addresses in the list of exclusions.

Adding networks to the excluded list negates the specified addresses and address blocks. That is, you can match any IP address with the exception of the excluded IP address or address blocks.

For example, excluding the literal address `192.168.1.1` specifies any IP address other than 192.168.1.1, and excluding `2001:db8:ca2e::fa4c` specifies any IP address other than 2001:db8:ca2e::fa4c.

You can exclude any combination of networks using literal or available networks. For example, excluding the literal values `192.168.1.1` and `192.168.1.5` *includes* any IP address other than 192.168.1.1 or 192.168.1.5. That is, the system interprets this as "**not** 192.168.1.1 **and not** 192.168.1.5," which matches any IP address other than those listed between brackets.

Note the following points when adding or editing network variables:

- You cannot logically exclude the value `any` which, if excluded, would indicate no address. For example, you cannot add a variable with the value `any` to the list of excluded networks.
- Network variables identify traffic for the specified intrusion rule and intrusion policy features. Note that preprocessor rules can trigger events regardless of the hosts defined by network variables used in intrusion rules.
- Excluded values must resolve to a subset of included values. For example, you cannot include the address block 192.168.5.0/24 and exclude 192.168.6.0/24.

Port Variables

Port variables represent TCP and UDP ports you can use in the **Source Port** and **Destination Port** header fields in intrusion rules that you enable in an intrusion policy. Port variables differ from port objects and port object groups in that port variables are specific to intrusion rules. You can create port objects for protocols other than TCP and UDP, and you can use port objects in various places in the system's web interface, including port variables, access control policies, network discovery rules, and event searches.

You can use port variables in the intrusion rule **Source Port** and **Destination Port** header fields to restrict packet inspection to packets originating from or destined to specific TCP or UDP ports.

When you use variables in these fields, the variable set you link to the intrusion policy associated with an access control rule or policy determines the values for these variables in the network traffic where you deploy the access control policy.

You can add any combination of the following port configurations to a variable:

- any combination of port variables and port objects that you select from the list of available ports

  Note that the list of available ports does not display port object groups, and you cannot add these to variables.

- individual port objects that you add from the New Variable or Edit Variable page, and can then add to your variable and to other existing and future variables

  Only TCP and UDP ports, including the value `any` for either type, are valid variable values. If you use the new or edit variables page to add a valid port object that is not a valid variable value, the object is added to the system but is not displayed in the list of available objects. When you use the object manager to edit a port object that is used in a variable, you can only change its value to a valid variable value.

- single, literal port values and port ranges

  You must separate port ranges with a dash (-). Port ranges indicated with a colon (:) are supported for backward compatibility, but you cannot use a colon in port variables that you create.

  You can list multiple literal port values and ranges by adding each individually in any combination.

Note the following points when adding or editing port variables:

- The default value for included ports in any variable you add is the word `any`, which indicates any port or port range. The default value for excluded ports is `none`, which indicates no ports.

Tip To create a variable with the value `any`, name and save the variable without adding a specific value.

- You cannot logically exclude the value `any` which, if excluded, would indicate no ports. For example, you cannot save a variable set when you add a variable with the value `any` to the list of excluded ports.
- Adding ports to the excluded list negates the specified ports and port ranges. That is, you can match any port with the exception of the excluded ports or port ranges.
- Excluded values must resolve to a subset of included values. For example, you cannot include the port range 10-50 and exclude port 60.

Advanced Variables

Advanced variables allow you to configure features that you cannot otherwise configure via the web interface. The system currently provides only one advanced variable, the USER_CONF variable.

USER_CONF

USER_CONF provides a general tool that allows you to configure one or more features not otherwise available via the web interface.

Caution Do **not** use the advanced variable USER_CONF to configure an intrusion policy feature unless you are instructed to do so in the feature description or by Support. Conflicting or duplicate configurations will halt the system.

When editing USER_CONF, you can type up to 4096 total characters on a single line; the line wraps automatically. You can include any number of valid instructions or lines until you reach the 8192 maximum character length for a variable or a physical limit such as disk space. Use the backslash (\) line continuation character after any complete argument in a command directive.

Resetting USER_CONF empties it.

Variable Reset

You can reset a variable to its default value on the variable set new or edit variables page. The following table summarizes the basic principles of resetting variables.

Table 77: Variable Reset Values

Resetting this variable type...	In this set type...	Resets it to...
default	default	the rule update value
user-defined	default	`any`
default or user-defined	custom	the current default set value (modified or unmodified)

Resetting a variable in a custom set simply resets it to the current value for that variable in the default set.

Conversely, resetting or modifying the value of a variable in the default set always updates the default value of that variable in all custom sets. When the reset icon is grayed out, indicating that you cannot reset the variable, this means that the variable has no customized value in that set. Unless you have customized the value for a variable in a custom set, a change to the variable in the default set updates the value used in any intrusion policy where you have linked the variable set.

Note It is good practice when you modify a variable in the default set to assess how the change affects any intrusion policy that uses the variable in a linked custom set, especially when you have not customized the variable value in the custom set.

You can hover your pointer over the **Reset icon** in a variable set to see the reset value. When the customized value and the reset value are the same, this indicates one of the following:

- you are in the custom or default set where you added the variable with the value `any`
- you are in the custom set where you added the variable with an explicit value and elected to use the configured value as the default value

Adding Variables to Sets

Adding a variable to a variable set adds it to all other sets. When you add a variable from a custom set, you must choose whether to use the configured value as the customized value in the default set:

- **If you use the configured value** (for example, 192.168.0.0/16), the variable is added to the default set using the configured value as a customized value with a default value of `any`. Because the current value in the default set determines the default value in other sets, the initial, default value in other custom sets is the configured value (which in the example is 192.168.0.0/16).
- **If you do not use the configured value**, the variable is added to the default set using only the default value `any` and, consequently, the initial, default value in other custom sets is `any`.

Example: Adding User-Defined Variables to Default Sets

The following diagram illustrates set interactions when you add the user-defined variable `Var1` to the default set with the value `192.168.1.0/24`.

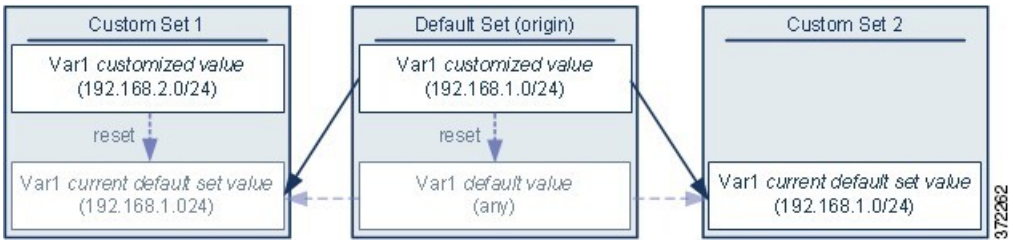

You can customize the value of `Var1` in any set. In Custom Set 2 where `Var1` has not been customized, its value is `192.168.1.0/24`. In Custom Set 1 the customized value `192.168.2.0/24` of `Var1` overrides the default value. Resetting a user-defined variable in the default set resets its default value to `any` in all sets.

It is important to note in this example that, if you do not update `Var1` in Custom Set 2, further customizing or resetting `Var1` in the default set consequently updates the current, default value of `Var1` in Custom Set 2, thereby affecting any intrusion policy linked to the variable set.

Although not shown in the example, note that interactions between sets are the same for user-defined variables and default variables except that resetting a default variable in the default set resets it to the value configured by Cisco in the current rule update.

Example: Adding User-Defined Variables to Custom Sets

The next two examples illustrate variable set interactions when you add a user-defined variable to a custom set. When you save the new variable, you are prompted whether to use the configured value as the default value for other sets. In the following example, you elect **to use** the configured value.

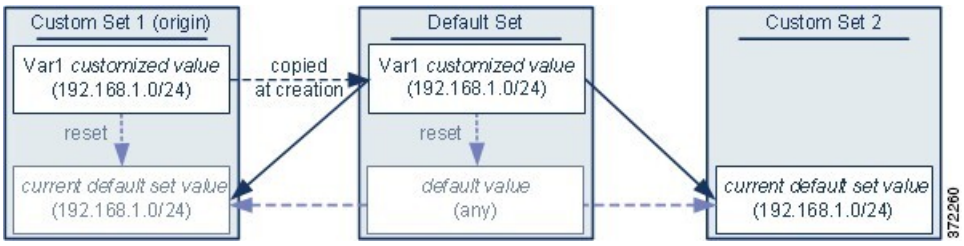

Note that, except for the origin of `Var1` from Custom Set 1, this example is identical to the example above where you added `Var1` to the default set. Adding the customized value `192.168.1.0/24` for `Var1` to Custom Set 1 copies the value to the default set as a customized value with a default value of `any`. Thereafter, `Var1` values and interactions are the same as if you had added `Var1` to the default set. As with the previous example, keep in mind that further customizing or resetting `Var1` in the default set consequently updates the current, default value of `Var1` in Custom Set 2, thereby affecting any intrusion policy linked to the variable set.

In the next example, you add `Var1` with the value 192.168.1.0/24 to Custom Set 1 as in the previous example, but you elect **not to use** the configured value of `Var1` as the default value in other sets.

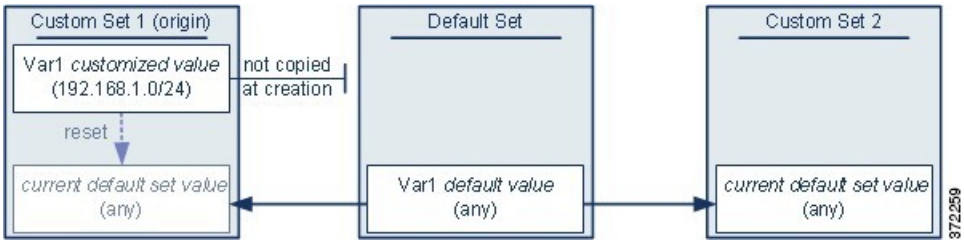

This approach adds `Var1` to all sets with a default value of `any`. After adding `Var1`, you can customize its value in any set. An advantage of this approach is that, by not initially customizing `Var1` in the default set, you decrease your risk of customizing the value in the default set and thus inadvertently changing the current value in a set such as Custom Set 2 where you have not customized `Var1`.

Nesting Variables

You can nest variables so long as the nesting is not circular. Nested, negated variables are not supported.

Valid Nested Variables

In this example, SMTP_SERVERS, HTTP_SERVERS, and OTHER_SERVERS are valid nested variables.

Variable	Type	Included Networks	Excluded Networks
SMTP_SERVERS	customized default	10.1.1.1	—
HTTP_SERVERS	customized default	10.1.1.2	—
OTHER_SERVERS	user-defined	10.2.2.0/24	—
HOME_NET	customized default	10.1.1.0/24 OTHER_SERVERS	SMTP_SERVERS HTTP_SERVERS

An Invalid Nested Variable

In this example, HOME_NET is an invalid nested variable because the nesting of HOME_NET is circular; that is, the definition of OTHER_SERVERS includes HOME_NET, so you would be nesting HOME_NET in itself.

Variable	Type	Included Networks	Excluded Networks
SMTP_SERVERS	customized default	10.1.1.1	—
HTTP_SERVERS	customized default	10.1.1.2	—
OTHER_SERVERS	user-defined	10.2.2.0/24 HOME_NET	—
HOME_NET	customized default	10.1.1.0/24 OTHER_SERVERS	SMTP_SERVERS HTTP_SERVERS

An Unsupported Nested, Negated Variable

Because nested, negated variables are not supported, you cannot use the variable NONCORE_NET as shown in this example to represent IP addresses that are outside of your protected networks.

Variable	Type	Included Networks	Excluded Networks
HOME_NET	customized default	10.1.0.0/16 10.2.0.0/16 10.3.0.0/16	—
EXTERNAL_NET	customized default	—	HOME_NET
DMZ_NET	user-defined	10.4.0.0/16	—
NOT_DMZ_NET	user-defined	—	DMZ_NET
NONCORE_NET	user-defined	EXTERNAL_NET NOT_DMZ_NET	—

Alternative to an Unsupported Nested, Negated Variable

As an alternative to the example above, you could represent IP addresses that are outside of your protected networks by creating the variable NONCORE_NET as shown in this example.

Variable	Type	Included Networks	Excluded Networks
HOME_NET	customized default	10.1.0.0/16 10.2.0.0/16 10.3.0.0/16	—
DMZ_NET	user-defined	10.4.0.0/16	—
NONCORE_NET	user-defined	—	HOME_NET DMZ_NET

Managing Variable Sets

To use variable sets, you must have the IPS license (for Firewall Threat Defense devices) or the Protection license (all other device types).

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Variable Set** from the list of object types.

Step 3 Manage your variable sets:

- Add — If you want to add a custom variable set, click **Add Variable Set**; see Creating Variable Sets, on page 1457.
- Delete — If you want to delete a custom variable set, click **Delete** ( ) next to the variable set, then click **Yes**. You cannot delete the default variable set or variable sets belonging to ancestor domains.

  Note
  Variables created in a variable set you delete are not deleted or otherwise affected in other sets.

- Edit — If you want to edit a variable set, click **Edit** ( ) next to the variable set you want to modify; see Editing Objects, on page 1357.
- Filter — If you want to filter variable sets by name, begin entering a name; as you type, the page refreshes to display matching names. If you want to clear name filtering, click **Clear** ( ) in the filter field.
- Manage Variables — To manage the variables included in variable sets, see Managing Variables, on page 1458.

Creating Variable Sets

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Variable Set** from the list of object types.

Step 3 Click **Add Variable Set**.

Step 4 Enter a **Name**.

Step 5 Optionally, enter a **Description**.

Step 6 Manage the variables in the set; see Managing Variables, on page 1458.

Step 7 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Managing Variables

You must have the IPS license (for Firewall Threat Defense devices) or the Protection license (all other device types).

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Variable Set** from the list of object types.

Step 3 Click **Edit** (✎) next to the variable set you want to edit.

If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Manage your variables:

- Display — If you want to display the complete value for a variable, hover your pointer over the value in the **Value** column next to the variable.
- Add — If you want to add a variable, click **Add**; see Adding Variables, on page 1459.
- Delete — Click **Delete** (🗑) next to the variable. If you have saved the variable set since adding the variable, click **Yes** to confirm that you want to delete the variable.

  You *cannot* delete the following:

  - default variables
  - user-defined variables that are used by intrusion rules or other variables
  - variables belonging to ancestor domains

- Edit — Click **Edit** (✎) next to the variable you want to edit; see Editing Variables, on page 1460.
- Reset — If you want to reset a modified variable to its default value, click **Reset** next to a modified variable. If reset is dimmed, one of the following is true:

  - The current value is already the default value.

- The configuration belongs to an ancestor domain.

Tip
Hover your pointer over an active reset to display the default value.

Step 5 Click **Save** to save the variable set. If the variable set is in use by an access control policy, click **Yes** to confirm that you want to save your changes.

Because the current value in the default set determines the default value in all other sets, modifying or resetting a variable in the default set changes the current value in other sets where you have not customized the default value.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Adding Variables

You must have the IPS license (for Firewall Threat Defense devices) or the Protection license (all other device types).

Procedure

Step 1 In the variable set editor, click **Add**.

Step 2 Enter a unique variable **Name**.

Step 3 From the **Type** drop-down list, choose either **Network** or **Port**.

Step 4 Specify values for the variable:

- If you want to move items from the list of available networks or ports to the list of included or excluded items, you can choose one or more items and then drag and drop, or click **Include** or **Exclude**.

  Tip
  If addresses or ports in the included and excluded lists for a network or port variable overlap, excluded addresses or ports take precedence.

- Enter a single literal value, then click **Add**. For network variables, you can enter a single IP address or address block. For port variables you can add a single port or port range, separating the upper and lower values with a hyphen (-). Repeat this step as needed to enter multiple literal values.
- If you want to remove an item from the included or excluded lists, click **Delete** ( ) next to the item.

Note
The list of items to include or exclude can be comprised of any combination of literal strings and existing variables, objects, and network object groups in the case of network variables.

Step 5 Click **Save** to save the variable. If you are adding a new variable from a custom set, you have the following options:

- Click **Yes** to add the variable using the configured value as the customized value in the default set and, consequently, the default value in other custom sets.
- Click **No** to add the variable as the default value of `any` in the default set and, consequently, in other custom sets.

Step 6 Click **Save** to save the variable set. Your changes are saved, and any access control policy the variable set is linked to displays an out-of-date status.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Editing Variables

You must have the IPS license (for Firewall Threat Defense devices) or the Protection license (all other device types).

You can edit both custom and default variables.

You cannot change the **Name** or **Type** values in an existing variable.

Procedure

Step 1 In the variable set editor, click **Edit** (✎) next to the variable you want to modify.

If **View** (👁) appears instead, the object belongs to an ancestor domain, or you do not have permission to modify the object.

Step 2 Modify the variable:

- If you want to move items from the list of available networks or ports to the list of included or excluded items, you can select one or more items and then drag and drop, or click **Include** or **Exclude**.

  Tip
  If addresses or ports in the included and excluded lists for a network or port variable overlap, excluded addresses or ports take precedence.

- Enter a single literal value, then click **Add**. For network variables, you can enter a single IP address or address block. For port variables you can add a single port or port range, separating the upper and lower values with a hyphen (-). Repeat this step as needed to enter multiple literal values.
- If you want to remove an item from the included or excluded lists, click **Delete** (🗑) next to the item.

Note
The list of items to include or exclude can be comprised of any combination of literal strings and existing variables, objects, and network object groups in the case of network variables.

Step 3 Click **Save** to save the variable.

Step 4 Click **Save** to save the variable set. If the variable set is in use by an access control policy, click **Yes** to confirm that you want to save your changes. Your changes are saved, and any access control policy the variable set is linked to displays an out-of-date status.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

VLAN Tag

Each VLAN tag object you configure represents a VLAN tag or range of tags.

You can group VLAN tag objects. Groups represent multiple objects; using a range of VLAN tags in a single object is not considered a group in this sense.

You can use VLAN tag objects and groups in various places in the system's web interface, including rules and event searches. For example, you could write an access control rule that applies only to a specific VLAN.

Creating VLAN Tag Objects

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **VLAN Tag** from the list of object types.

Step 3 Choose **Add Object** from the **Add VLAN Tag** drop-down list.

Step 4 Enter a **Name**.

Step 5 Enter a **Description**.

Step 6 Enter a value in the **VLAN Tag** field. Use a hyphen to specify a range of VLAN tags.

Step 7 Manage overrides for the object:

- If you want to allow overrides for this object, check the **Allow Overrides** check box; see Allowing Object Overrides, on page 1362.
- If you want to add override values to this object, expand the Override section and click **Add**; see Adding Object Overrides, on page 1362.

Step 8 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

VPN

You can use the following VPN objects on Firewall Threat Defense devices. To use these objects, you must have Admin privileges, and your Smart License account must satisfy export controls. You can configure these objects in leaf domains only.

Certificate Map Objects

Certificate Map objects are a named set of certificate matching rules. These objects are used to provide an association between a received certificate and a Remote Access VPN connection profile. Connection Profiles and Certificate Map objects are both part of a remote access VPN policy. If a received certificate matches the rules contained in the certificate map, the connection is "mapped", or associated with the specified connection profile. The rules are in priority order, they are matched in the order they are shown in the UI. The matching ends when the first rule within the Certificate Map object results in a match.

Navigation

Objects > **Object Management** > **VPN** > **Certificate Map**

Fields

- **Name**—Identify this object so it can be referred to from other configurations, such as Remote Access VPN.
- **Mapping Criteria**—Specify the contents of the certificate to evaluate. If the certificate satisfies these rules, the user will be mapped to the connection profile containing this object.
  - **Field**—Select the field for the matching rule according to the Subject or the Issuer of the client certificate.

    If the **Field** is set to *Alternative Subject* or *Extended Key Usage* the Component will be frozen as *Whole Field*
  - **Component**—Select the component of the client certificate to use for the matching rule.

Note **SER (Serial Number) component** - Ensure you specify the serial number for the Subject field. The certificate map only matches with a serial number attribute in the subject name.

  - **Operator**—Select the operator for the matching rule as follows:
    - Equals—The certificate component must match the entered value. If they do not match exactly, the connection is denied.
    - Contains—The certificate component must contain the entered value. If the component does not contain the value, the connection is denied.
    - Does Not Equal—The certificate component cannot equal the entered value. For example, for a selected certificate component of Country, and an entered value of US, if the client county value equals US, then the connection is denied.

- Does Not Contain—The certificate component cannot contain the entered value. For example, for a selected certificate component of Country, and an entered value of US, if the client county value contains US, the connection is denied.

- **Value**—The value of the matching rule. The value entered is associated with the selected component and operator.

Related Topics

Configure Certificate Maps, on page 1618

Secure Client Custom Attributes Objects

Custom attributes are used by the Secure Client to configure features such as Per App VPN, Allow or defer upgrade, and Dynamic split tunneling. A custom attribute has a type and a named value. The type of the attribute is defined first, then one or more named values of this type can be defined. You can create the Secure Client custom attributes objects using the Firewall Management Center, add the objects to a group policy and associate the group policy with a remote access VPN to enable the features for the VPN clients.

Firewall Threat Defense supports the following features using the custom attribute objects:

- **Per App VPN**—The Per App VPN feature helps identify an app and tunnel only applications allowed by the Firewall Threat Defense administrator over the VPN.
- **Allow or defer upgrade**— Deferred Upgrade allows the Secure Client user to delay download of the Secure Client upgrade. When a client update is available, you can configure the attributes for Secure Client to open a dialog asking the user if they would like to update, or to defer the upgrade.
- **Dynamic Split Tunneling**— With dynamic split tunneling, you can provision policies that either include or exclude IP addresses or networks from the VPN tunnel. Dynamic split tunneling is configured by creating a custom attribute and adding it to a group policy.

For step-by-step instructions to configure Secure Client custom attributes, see Add Secure Client Custom Attributes Objects, on page 1463 and

For details about the specific custom attributes to configure for a feature, see the *Cisco Secure Client (including AnyConnect) Administrator Guide* for the Secure Client release you are using.

Related Topics

Group Policy Secure Client Options, on page 1469

Add Secure Client Custom Attributes Objects

Before you begin

Ensure that you have done the following before adding a custom attribute object for Per App VPN:

- Per App VPN must be properly configured via MDM and each device must be enrolled to the MDM server
- Create a base64 encoded string for each app using the Cisco Secure Client Enterprise Application Selector Tool.

1. Download the Cisco Secure Client Enterprise Application Selector Tool from here.

2. Open the Application Selection Tool and select the mobile platform from the drop-down menu located on the upper left.
3. Add rule by entering Friendly name and App ID; rest of the fields are optional.
4. On the menu bar, click on **Policy**. The encoded base65 rule is displayed in its encoded format.
5. Select and copy the policy string, and save it to use later when you create the Secure Client Custom Attributes object.

Procedure

Step 1 Choose **Objects** > **Object Management** > **VPN** > **Custom Attribute**.

Step 2 Click **Secure Client Custom Attribute**.

Step 3 Enter a **Name** and optionally a **Description** for the attribute.

Step 4 Select an attribute from the **Secure Client Attribute** drop-down list:

- **Per App VPN** — Select this option and specify the base64 encoded string in the **Attribute Value** box.
- **Allow Defer Update**—Select one of the following options and specify the required information to allow or defer Secure Client update:
  - **Show the prompt until user takes action**—Display the prompt to the VPN user until the user chooses to allow or defer the VPN client update.
  - **Show the prompt until times out**—Choose this option to display the prompt for a specified duration and specify the duration int the **Timeout** box.
  - **Do not show the prompt and take automatic action**—Choose this option to automatically allow or defer the VPN update.
  - **Default Action**—Select the default action to be taken when the user does not respond, or when you want to configure an automatic action without the user's intervention. You can choose to update the Secure Client or postpone the update.
  - **Minimum Version**—Specify the minimum Secure Client version to be present on the client system to allow or defer the update.
- **Dynamic Split Tunneling**—Select this option to include or exclude IP addresses or networks from the VPN tunnel.
  - **Include domains**—Specify domain names that will be included in the remote access VPN tunnel.
  - **Exclude domains**—Specify domain names that will be excluded from the remote access VPN tunnel.

Step 5 Select the **Allow Overrides** check box to allow object overrides.

Step 6 Click **Save**.
The custom attributes object is added to the list.

What to do next

Associate the custom attributes with a group policy. See Add Custom Attributes to a Group Policy, on page 1465 .

Add Custom Attributes to a Group Policy

You must associate Secure Client custom attributes with a group policy to use them for remote access VPN connections. You

Procedure

Step 1 Select **Objects** > **Object Management** > **VPN** > **Group Policy**.

Step 2 Add a new group policy or edit an existing group policy.

Step 3 Click **Secure Client** > **Custom Attributes**.

Step 4 Click **Add**.

Step 5 Select the **Secure Client Attribute**: Per App VPN, Allow Defer Update, or Dynamic Split Tunneling.

Step 6 Select a **Custom Attribute Object** from the list.

Note
Click Add (+) to create a new custom attribute object for the selected Secure Client attribute. You can also create a custom attribute object at **Objects** > **Object Management** > **VPN** > **Custom Attribute**. See Add Secure Client Custom Attributes Objects, on page 1463.

Step 7 Click **Add** to save the attributes to the group policy and then click **Save** to save the changes to the group policy.

Related Topics

Group Policy Secure Client Options, on page 1469

Firewall Threat Defense Group Policy Objects

A group policy is a set of attribute and value pairs, stored in a group policy object, that define the remote access VPN experience. For example, in the group policy object, you configure general attributes such as addresses, protocols, and connection settings.

The group policy applied to a user is determined when the VPN tunnel is being established. The RADIUS authorization server assigns the group policy, or it is obtained from the current connection profile.

Note There is no group policy attribute inheritance on the Firewall Threat Defense. A group policy object is used, in its entirety, for a user. The group policy object identified by the AAA server upon login is used, or, if that is not specified, the default group policy configured for the VPN connection is used. The provided default group policy can be set to your default values, but will only be used if it is assigned to a connection profile and no other group policy has been identified for the user.

To use group objects, you must have one of these Secure Client licenses associated with your Smart License account with Export-Controlled Features enabled:

- Secure Client VPN Only
- Secure Client Advantage
- Secure Client Premier

Related Topics

Configure Group Policy Objects, on page 1466

Configure Group Policy Objects

See Firewall Threat Defense Group Policy Objects, on page 1465.

Procedure

Step 1 Choose **Objects** > **Object Management** > **VPN** > **Group Policy**.

Previously configured policies are listed including the system default. Depending on your level of access, you may edit, view, or delete a group policy.

Step 2 Click **Add Group Policy** or choose a current policy to edit.

Step 3 Enter a **Name** and optionally a **Description** for this policy.

The name can be up to 64 characters, spaces are allowed. The description can be up to 1,024 characters.

Step 4 Specify the **General** parameters for this Group Policy as described in Group Policy General Options, on page 1466.

Step 5 Specify the **Secure Client** parameters for this Group Policy as described in Group Policy Secure Client Options, on page 1469.

Step 6 Specify the **Advanced** parameters for this Group Policy as described in Group Policy Advanced Options, on page 1472.

Step 7 Click **Save**.
The new Group Policy is added to the list.

What to do next

Add the group policy object to a remote access VPN connection profile.

Group Policy General Options

Navigation Path

Objects > **Object Management** > **VPN** > **Group Policy**, Click **Add Group Policy** or choose a current policy to edit., then select the **General** tab.

VPN Protocols Fields

Specify the types of Remote Access VPN tunnels that can be used when applying this group policy. **SSL** or **IPsec IKEv2**.

IP Address Pools

Specifies the IPv4 address assignment that is applied based on address pools that are specific to user-groups in Remote Access VPN. For Remote Access VPN, you can assign IP address from specific address pools for identified user groups using RADIUS/ISE for authorization. You can seamlessly perform policy enforcement for user or user groups in systems which are not identity-aware, by configuring particular Group Policy as RADIUS Authorization attribute (GroupPolicy/Class), for a particular user group. For example, you have to select a specific address pool for contractors and policy enforcement using those addresses to allow restricted access to internal network.

The order of preference that Firewall Threat Defense device assigns the IPv4 Address Pools to the clients:

1. RADIUS attribute for IPv4Address Pool
2. RADIUS attribute for Group Policy
3. Address Pool in Group Policy mapped to a Connection Profile
4. IPv4Address Pool in Connection Profile

Some limitations around using IP address pools in Group Policy:

- IPv6 address pool is not supported.
- Maximum of six IPv4 address pools can be configured in a Group Policy.
- Deployment failures are seen when address pools in use are modified. You must logoff all the users before making any changes to the address pools.
- When address pools are renamed or overlapping address pools are configured, deployment could fail. You must deploy the changes by removing the old address pool and later deploying the changed address pool.

  Some troubleshooting commands :

  - `show ip local pool <address-pool-name>`
  - `show vpn-sessiondb detail anyconnect`
  - `vpn-sessiondb loggoff all noconfirm`

Banner Fields

Specifies the banner text to present to users at login. The length can be up to 491 characters. There is no default value. The IPsec VPN client supports full HTML for the banner, however, the Secure Client supports only partial HTML. To ensure that the banner displays properly to remote users, use the /n tag for IPsec clients, and the <BR> tag for SSL clients.

DNS/WINS Fields

Domain Naming System (DNS) and Windows Internet Naming System (WINS) servers. Used for Secure Client name resolution.

- **Primary DNS Server** and **Secondary DNS Server**—Choose or create a Network Object which defines the IPv4 or IPv6 addresses of the DNS servers you want this group to use.
- **Primary WINS Server** and **Secondary WINS Server**—Choose or create a Network Object containing the IP addresses of the WINS servers you want this group to use.

- **DHCP Network Scope**—Choose or create a Network Object containing a routable IPv4 address on the same subnet as the desired pool, but not within the pool. The DHCP server determines which subnet this IP address belongs to and assigns an IP address from that pool. If not set properly, deployment of the VPN policy fails.

  If you configure DHCP servers for the address pool in the connection profile, the DHCP scope identifies the subnets to use for the pool for this group. The DHCP server must also have addresses in the same subnet identified by the scope. The scope allows you to select a subset of the address pools defined in the DHCP server to use for this specific group.

  If you do not define a network scope, the DHCP server assigns IP addresses in the order of the address pools configured. It goes through the pools until it identifies an unassigned address.

  We recommend using the IP address of an interface whenever possible for routing purposes. For example, if the pool is 10.100.10.2-10.100.10.254, and the interface address is 10.100.10.1/24, use 10.100.10.1 as the DHCP scope. Do not use the network number. You can use DHCP for IPv4 addressing only. If the address you choose is not an interface address, you might need to create a static route for the scope address.

  LINK-SELECTION (RFC 3527) and SUBNET-SELECTION (RFC 3011) are currently not supported.

- **Default Domain**—Name of the default domain. Specify a top-level domain, for example, example.com.

Split Tunneling Fields

Split tunneling directs some network traffic through the VPN tunnel (encrypted) and the remaining network traffic outside the VPN tunnel (unencrypted or "in the clear").

- **IPv4 Split Tunneling / IPv6 Split Tunneling**—By default, split tunneling is not enabled. For both IPv4 and IPv6 it is set to **Allow all traffic over tunnel**. Left as is, all traffic from the endpoint goes over the VPN connection.

  To configure split tunneling, choose the **Tunnel networks specified below** or **Exclude networks specified below** policy. Then configure an access control list for that policy.

- **Split Tunnel Network List Type**—Choose the type of Access List you are using. Then choose or create a **Standard Access List** or **Extended Access List**. See Access List, on page 1369 for details.

Note From Version 7.4, the **Standard Access List** option is disabled by default. Use the **Extended Access List** option to add an access list.

- **DNS Request Split Tunneling**—Also known as Split DNS. Configure the DNS behavior expected in your environment.

  By default, split DNS is not enabled and set to **Send DNS request as per split tunnel policy**. Choosing **Always send DNS request over tunnel** forces all DNS requests to be sent over the tunnel to the private network.

  To configure split DNS, choose **Send only specified domains over tunnel**, and enter the list of domain names in the **Domain List** field. These requests are resolved through the split tunnel to the private network. All other names are resolved using the public DNS server. Enter up to ten entries in the list of domains, separated by commas. The entire string can be no longer than 491 characters.

Related Topics

Configure Group Policy Objects, on page 1466

Group Policy Secure Client Options

These specifications apply to the operation of the Secure Client VPN.

Navigation

Objects > **Object Management** > **VPN** > **Group Policy**. Click **Add Group Policy** or choose a current policy to edit. Then select the **Secure Client** tab.

Profile Fields

Profile—Choose or create a file object containing the Secure Client Profile. See File Objects, on page 1478 for object creation details.

The Secure Client Profile is a group of configuration parameters stored in an XML file. The Secure Client software uses it to configure the connection entries that appear in the client's user interface. These parameters (XML tags) also configure settings to enable more Secure Client features.

Use the GUI-based Secure Client Profile Editor, an independent configuration tool, to create the Secure Client Profile. See the *Secure Client Profile Editor* chapter in the appropriate release of the Cisco Secure Client (including AnyConnect) Administrator Guide for details.

Management Profile Fields

A Management VPN Tunnel provides connectivity to the corporate network whenever the endpoint is powered up, even if end-user does not connect over VPN.

Management VPN Profile—The Management Profile file contains settings for enabling and establishing Management VPN Tunnel on endpoint.

The Standalone Management VPN Tunnel profile editor can be used to create a new profile file or modify an existing file. You can download the profile editor from Cisco Software Download Center.

For more information about adding a profile file, see File Objects, on page 1478.

Client Modules Fields

Cisco Secure Client VPN Only offers enhanced security through various built-in modules. These modules provide services such as web security, network visibility into endpoint flows, and off-network roaming protection. Each client module includes a client profile that includes a group of custom configurations as per your requirement.

The following Secure Client modules are optional and you can configure these modules to be downloaded when a VPN user downloads Secure Client:

- **AMP Enabler**—Deploys advanced malware protection (AMP) for endpoints.
- **DART**—Captures a snapshot of system logs and other diagnostic information, which can be sent to the Cisco TAC for troubleshooting.
- **ISE Posture**—Uses the OPSWAT library to perform posture checks to assess an endpoint's compliance.
- **Network Access Manager**—Provides 802.1X (Layer 2) and device authentication for access to both wired and wireless networks.

- **Network Visibility**—Enhances the enterprise administrator's ability to do capacity and service planning, auditing, compliance, and security analytics.
- **Start Before Login**—Forces the user to connect to the enterprise infrastructure over a VPN connection before logging on to Windows by starting Secure Client before the Windows login dialog box appears.
- **Umbrella Roaming Security**—Provides DNS-layer security when no VPN is active.
- **Web Security**—Analyzes the elements of a web page, allows acceptable content, and blocks malicious or unacceptable content based on a defined security policy.

Click **Add** and select the following for each client module:

- **Client Module**—Select the Secure Client module from the list.
- **Profile to download**—Choose or create a file object containing the Secure Client Profile. See File Objects, on page 1478 for object creation details.
- **Enable module download**—Select to enable endpoints to download the client module along with the profile. If not selected, the endpoints can download only the client profile.

Use the GUI-based Secure Client Profile Editor, an independent configuration tool to create a client profile for each module. Download the Secure Client Profile Editor from Cisco Software Download Center. See the *Secure Client Profile Editor* chapter in the appropriate release of the Cisco Secure Client (including AnyConnect) for details.

SSL Settings Fields

- **SSL Compression**—Whether to enable data compression, and if so, the method of data compression to use, Deflate, or LZS. SSL Compression is Disabled by default.

  Data compression speeds up transmission rates, but also increases the memory requirement and CPU usage for each user session. Therefore, decreasing the overall throughput of the security appliance.
- **DTLS Compression**—Whether to compress Datagram Transport Layer Security (DTLS) connections for this group using LZS or not. DTLS Compression is Disabled by default.
- **MTU Size**—The maximum transmission unit (MTU) size for SSL VPN connections established by the Cisco Secure Client VPN Only. Default is 1406 Bytes, valid range is 576 to 1462 Bytes.
  - **Ignore DF Bit**—Whether to ignore the Don't Fragment (DF) bit in packets that need fragmentation. Allows the forced fragmentation of packets that have the DF bit set, allowing them to pass through the tunnel.

Connection Settings Fields

- **Enable Keepalive Messages between Secure Client and VPN gateway**. And its **Interval** setting.—Whether to exchange keepalive messages between peers to demonstrate that they are available to send and receive data in the tunnel. Default is enabled. Keepalive messages transmit at set intervals. If enabled, enter the time interval (in seconds) that the remote client waits between sending IKE keepalive packets. The default interval is 20 seconds, the valid range is 15 to 600 seconds.
- **Enable Dead Peer Detection on ...**. And their **Interval** settings.—Dead Peer Detection (DPD) ensures that the VPN secure gateway or the VPN client quickly detects when the peer is no longer responding, and the connection has failed. Default is enabled for both the gateway and the client. DPD messages

transmit at set intervals. If enabled, enter the time interval (in seconds) that the remote client waits between sending DPD messages. The default interval is 30 seconds, the valid range is 5 to 3600 seconds.

- **Enable Client Bypass Protocol**—Allows you to configure how the secure gateway manages IPv4 traffic (when it is expecting only IPv6 traffic), or how it manages IPv6 traffic (when it is expecting only IPv4 traffic).

  When the Secure Client makes a VPN connection to the headend, the headend assigns it an IPv4, IPv6, or both an IPv4 and IPv6 address. If the headend assigns the Secure Client connection only an IPv4 address or only an IPv6 address, you can configure the Client Bypass Protocol to drop network traffic for which the headend did not assign an IP address (default, disabled, not checked), or allow that traffic to bypass the headend and be sent from the client unencrypted or "in the clear" (enabled, checked).

  For example, assume that the secure gateway assigns only an IPv4 address to the Secure Client connection and the endpoint is dual-stacked. When the endpoint attempts to reach an IPv6 address, if Client Bypass Protocol is disabled, the IPv6 traffic is dropped; however, if Client Bypass Protocol is enabled, the IPv6 traffic is sent from the client in the clear.

- **SSL rekey**—Enables the client to rekey the connection, renegotiating the crypto keys and initialization vectors, increasing the security of the connection. This is disabled by default. When enabled, the renegotiation can be done at a specified interval and rekey the existing tunnel or create a new tunnel by setting the following fields:
  - **Method**—Available when SSL rekey is enabled. Create a **New Tunnel** (default), or renegotiate, the **Existing Tunnel**'s specifications.
  - **Interval**—Available when SSL rekey is enabled. Set to a default of 4 minutes with a range of 4-10080 minutes (1 week).

- **Client Firewall Rules**—Use the Client Firewall Rules to configure firewall settings for the VPN client's platform. Rules are based on criteria such as source address, destination address, and protocol. Extended Access Control List building block objects are used to define the traffic filter criteria. Choose or create an Extended ACL for this group policy. Define a **Private Network Rule** to control data flowing to the private network, a **Public Network Rule** to control data flowing "in the clear", outside of the established VPN tunnel, or both.

Note Ensure that the ACL contains only TCP/UDP/ICMP/IP ports and source network as any, any-ipv4 or any-ipv6.

Only VPN clients running Microsoft Windows can use these firewall settings.

Custom Attributes Fields

This section lists the Secure Client Custom attributes that are used by the Secure Client to configure features such as Per App VPN, Allow or defer upgrade, and Dynamic split tunneling. Click **Add** to add custom attributes to the group policy.

1. Select the **Secure Client Attrinute**: Per App VPN, Allow Defer Update, or Dynamic Split Tunneling.
2. Select a **Custom Attribute Object** from the list.

Note Click Add (+) to create a new custom attribute object for the selected Secure Client attribute. You can also create a custom attribute object at **Objects** > **Object Management** > **VPN** > **Custom Attribute**. See Add Secure Client Custom Attributes Objects, on page 1463.

3. Click **Add** to save the attributes to the group policy and then click **Save** to save the changes to the group policy.

Related Topics

Configure Group Policy Objects, on page 1466

Group Policy Advanced Options

Navigation Path

Objects > **Object Management** > **VPN** > **Group Policy**, Click **Add Group Policy** or choose a current policy to edit., then select the **Advanced** tab.

Traffic Filter Fields

- **Access List Filter**—Filters consist of rules that determine whether to allow or block tunneled data packets coming through the VPN connection. Rules are based on criteria such as source address, destination address, and protocol. Note that the VPN filter applies to initial connections only. It does not apply to secondary connections, such as a SIP media connection, that are opened due to the action of application inspection. Extended Access Control List building block objects are used to define the traffic filter criteria. Choose or create a new Extended ACL for this group policy.
- **Restrict VPN to VLAN**—Also called "VLAN mapping," this parameter specifies the egress VLAN interface for sessions to which this group policy applies. The ASA forwards all traffic from this group to the selected VLAN.

  Use this attribute to assign a VLAN to the group policy to simplify access control. Assigning a value to this attribute is an alternative to using ACLs to filter traffic on a session. In addition to the default value (Unrestricted), the drop-down list shows only the VLANs that are configured in this ASA. Allowed values range from 1 to 4094.

Session Settings Fields

- **Access Hours**—Choose or create a time range object. This object specifies the range of time this group policy is available to be applied to a remote access user. See Time Range, on page 1442 for details.
- **Simultaneous Logins Per User**—Specifies the maximum number of simultaneous logins allowed for a user. The default value is 3. The minimum value is 0, which disables login and prevents user access. Allowing several simultaneous connections may compromise security and affect performance.
- **Maximum Connection Time / Alert Interval**—Specifies the maximum user connection time in minutes. At the end of this time, the system stops the connection. The minimum is 1 minute). The Alert interval specifies the interval of time before maximum connection time is reached to display a message to the user.
- **Idle Timeout / Alert Interval**—Specifies this user's idle timeout period in minutes. If there is no communication activity on the user connection in this period, the system stops the connection. The

minimum time is 1 minute. The default is 30 minutes. The Alert interval specifies the interval of time before idle time is reached to display a message to the user.

Related Topics

Configure Group Policy Objects, on page 1466

Firewall Threat Defense IPsec Proposals

IPsec Proposals (or Transform Sets) are used when configuring VPN topologies. During the IPsec security association negotiation with ISAKMP, the peers agree to use a particular proposal to protect a particular data flow. The proposal must be the same for both peers.

There are separate IPsec proposal objects based on the IKE version, IKEv1, or IKEv2:

- When you create an IKEv1 IPsec Proposal (Transform Set) object, you select the mode in which IPsec operates, and define the required encryption and authentication types. You can select single options for the algorithms. If you want to support multiple combinations in a VPN, create multiple IKEv1 IPsec Proposal objects.
- When you create an IKEv2 IPsec Proposal object, you can select all of the encryption and Hash Algorithms allowed in a VPN. During IKEv2 negotiations, the peers select the most appropriate options that each support.

The Encapsulating Security Protocol (ESP) is used for both IKEv1 and IKEv2 IPsec Proposals. It provides authentication, encryption, and antireplay services. ESP is IP protocol type 50.

Note We recommend using both encryption and authentication on IPsec tunnels.

Configure IKEv1 IPsec Proposal Objects

Procedure

Step 1 Choose **Objects** > **Object Management** and then **VPN** > **IPsec IKev1 Proposal** from the table of contents.

Previously configured Proposals are listed including system defined defaults. Depending on your level of access, you may **Edit** (), **View** (), or **Delete** () a Proposal.

Step 2 Choose **Add IPsec IKEv1 Proposal** to create a new Proposal.

Step 3 Enter a **Name** for this Proposal

The name of the policy object. A maximum of 128 characters is allowed.

Step 4 Enter a **Description** for this Proposal.

A description of the policy object. A maximum of 1024 characters is allowed.

Step 5 Choose the **ESP Encryption** method. The Encapsulating Security Protocol (ESP) encryption algorithm for this Proposal.

For IKEv1, select one of the options. When deciding which encryption and Hash Algorithms to use for the IPsec proposal, your choice is limited to algorithms supported by the devices in the VPN. For a full explanation of the options, see Deciding Which Encryption Algorithm to Use, on page 1502.

Step 6 Select an option for **ESP Hash**.

For a full explanation of the options, see Deciding Which Hash Algorithms to Use, on page 1503.

Step 7 Click **Save**
The new Proposal is added to the list.

Configure IKEv2 IPsec Proposal Objects

Procedure

Step 1 Choose **Objects** > **Object Management** and then **VPN** > **IKEv2 IPsec Proposal** from the table of contents.

Previously configured Proposals are listed including system defined defaults. Depending on your level of access, you may **Edit** (), **View** (), or **Delete** () a Proposal.

Step 2 Choose **Add IKEv2 IPsec Proposal** to create a new proposal.

Step 3 Enter a **Name** for this proposal.

The name of the policy object. A maximum of 128 characters is allowed.

Step 4 Enter a **Description** for this proposal.

A description of the policy object. A maximum of 1024 characters is allowed.

Step 5 Choose the **ESP Hash** method, the hash or integrity algorithm to use in the proposal for authentication.

Note
Firewall Threat Defense does not support IPSec tunnels with NULL encryption. Make sure that you do not choose NULL encryption for IPSec IKEv2 proposal.

For IKEv2, select all the options you want to support for **ESP Hash**. For a full explanation of the options, see Deciding Which Hash Algorithms to Use, on page 1503.

Step 6 Choose the **ESP Encryption** method. The Encapsulating Security Protocol (ESP) encryption algorithm for this proposal.

For IKEv2, click Select to open a dialog box where you can select all of the options you want to support. When deciding which encryption and Hash Algorithms to use for the IPsec proposal, your choice is limited to algorithms supported by the devices in the VPN. For a full explanation of the options, see Deciding Which Encryption Algorithm to Use, on page 1502.

Step 7 Click **Save**
The new proposal is added to the list.

Firewall Threat Defense IKE Policies

Internet Key Exchange (IKE) is a key management protocol that is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and automatically establish IPsec security associations (SAs). The IKE negotiation comprises two phases. Phase 1 negotiates a security association between two IKE peers, which enables the peers to communicate securely in Phase 2. During Phase 2 negotiation, IKE establishes SAs for other applications, such as IPsec. Both phases use proposals when they negotiate a connection. An IKE proposal is a set of algorithms that two peers use to secure the negotiation between them. IKE negotiation begins by each peer agreeing on a common (shared) IKE policy. This policy states which security parameters are used to protect subsequent IKE negotiations.

For IKEv1, IKE proposals contain a single set of algorithms and a modulus group. You can create multiple, prioritized policies to ensure that at least one policy matches a remote peer's policy. Unlike IKEv1, in an IKEv2 proposal, you can select multiple algorithms and modulus groups in one policy. Since peers choose during the Phase 1 negotiation, this makes it possible to create a single IKE proposal, but consider multiple, different proposals to give higher priority to your most desired options. For IKEv2, the policy object does not specify authentication, other policies must define the authentication requirements.

An IKE policy is required when you configure a site-to-site IPsec VPN. For more information, see VPN, on page 1495.

Configure IKEv1 Policy Objects

Use the IKEv1 Policy page to create, delete, or edit an IKEv1 policy object. These policy objects contain the parameters required for IKEv1 policies.

Procedure

Step 1 Choose **Objects** > **Object Management** and then **VPN** > **IKEv1 Policy** from the table of contents.

Previously configured policies are listed including system defined defaults. Depending on your level of access, you may **Edit** (), **View** (), or **Delete** () a proposal.

Step 2 (Optional) Choose **Add IKEv1 Policy** to create a new policy object.

Step 3 Enter a **Name** for this policy. A maximum of 128 characters is allowed.

Step 4 (Optional) Enter a **Description** for this proposal. A maximum of 1,024 characters is allowed.

Step 5 Enter the **Priority** value of the IKE policy.

The priority value determines the order of the IKE policy compared by the two negotiating peers when attempting to find a common security association (SA). If the remote IPsec peer does not support the parameters selected in your first priority policy, it tries to use the parameters defined in the next lowest priority. Valid values range from 1 to 65,535. The lower the number, the higher the priority. If you leave this field blank, Management Center assigns the lowest unassigned value starting with 1, then 5, then continuing in increments of 5.

Step 6 Choose the **Encryption** method.

When deciding which encryption and Hash Algorithms to use for the IKEv1 policy, your choice is limited to algorithms supported by the peer devices. For an extranet device in the VPN topology, you must choose the algorithm that matches both peers. For IKEv1, select one of the options. For a full explanation of the options, see Deciding Which Encryption Algorithm to Use, on page 1502.

Step 7 Choose the **Hash** Algorithm that creates a Message Digest, which is used to ensure message integrity.

When deciding which encryption and Hash Algorithms to use for the IKEv1 proposal, your choice is limited to algorithms supported by the managed devices. For an extranet device in the VPN topology, you must choose the algorithm that matches both peers. For a full explanation of the options, see Deciding Which Hash Algorithms to Use, on page 1503.

Step 8 Set the **Diffie-Hellman Group**.

The Diffie-Hellman group to use for encryption. A larger modulus provides higher security but requires more processing time. The two peers must have a matching modulus group. Select the group that you want to allow in the VPN. For a full explanation of the options, see Deciding Which Diffie-Hellman Modulus Group to Use, on page 1503.

Step 9 Set the **Lifetime** of the security association (SA), in seconds. You can specify a value from 120 to 2,147,483,647 seconds. The default is 86400.

When the lifetime is exceeded, the SA expires and must be renegotiated between the two peers. Generally, the shorter the lifetime (up to a point), the more secure your IKE negotiations. However, with longer lifetimes, future IPsec security associations can be set up more quickly than with shorter lifetimes.

Step 10 Set the **Authentication Method** to use between the two peers.

- **Preshared Key**—Preshared keys allow for a secret key to be shared between two peers and to be used by IKE during the authentication phase. If one of the participating peers is not configured with the same preshared key, the IKE SA cannot be established.
- **Certificate**—When you use Certificates as the authentication method for VPN connections, peers obtain digital certificates from a CA server in your PKI infrastructure, and trade them to authenticate each other.

Note
In a VPN topology that supports IKEv1, the **Authentication Method** specified in the chosen IKEv1 Policy object becomes the default in the IKEv1 **Authentication Type** setting. These values must match, otherwise, your configuration will error.

Step 11 Click **Save**
The new IKEv1 policy is added to the list.

Configure IKEv2 Policy Objects

Use the IKEv2 policy dialog box to create, delete, and edit an IKEv2 policy object. These policy objects contain the parameters required for IKEv2 policies.

Procedure

Step 1 Choose **Objects** > **Object Management** and then **VPN** > **IKEv2 Policy** from the table of contents.

Previously configured policies are listed including system defined defaults. Depending on your level of access, you may **Edit** (), **View** (), or **Delete** () a policy.

Step 2 Choose **Add IKEv2 Policy** to create a new policy.

Step 3 Enter a **Name** for this policy.

The name of the policy object. A maximum of 128 characters is allowed.

Step 4 Enter a **Description** for this policy.

A description of the policy object. A maximum of 1024 characters is allowed.

Step 5 Enter the **Priority**.

The priority value of the IKE proposal. The priority value determines the order of the IKE proposals compared by the two negotiating peers when attempting to find a common security association (SA). If the remote IPsec peer does not support the parameters selected in your first priority policy, it tries to use the parameters defined in the next lowest priority policy. Valid values range from 1 to 65535. The lower the number, the higher the priority. If you leave this field blank, Management Center assigns the lowest unassigned value starting with 1, then 5, then continuing in increments of 5.

Step 6 Set the **Lifetime** of the security association (SA), in seconds. You can specify a value from 120 to 2,147,483,647 seconds. The default is 86400.

When the lifetime is exceeded, the SA expires and must be renegotiated between the two peers. Generally, the shorter the lifetime (up to a point), the more secure your IKE negotiations. However, with longer lifetimes, future IPsec security associations can be set up more quickly than with shorter lifetimes.

Step 7 Choose the **Integrity Algorithms** portion of the Hash Algorithm used in the IKE policy. The Hash Algorithm creates a Message Digest, which is used to ensure message integrity.

When deciding which encryption and Hash Algorithms to use for the IKEv2 proposal, your choice is limited to algorithms supported by the managed devices. For an extranet device in the VPN topology, you must choose the algorithm that matches both peers. Select all the algorithms that you want to allow in the VPN.For a full explanation of the options, see Deciding Which Hash Algorithms to Use, on page 1503.

Step 8 Choose the **Encryption Algorithm** used to establish the Phase 1 SA for protecting Phase 2 negotiations.

When deciding which encryption and Hash Algorithms to use for the IKEv2 proposal, your choice is limited to algorithms supported by the managed devices. For an extranet device in the VPN topology, you must choose the algorithm that matches both peers. Select all the algorithms that you want to allow in the VPN. For a full explanation of the options, see Deciding Which Encryption Algorithm to Use, on page 1502.

Step 9 Choose the **PRF Algorithm**.

The pseudorandom function (PRF) portion of the Hash Algorithm used in the IKE policy. In IKEv1, the Integrity and PRF algorithms are not separated, but in IKEv2, you can specify different algorithms for these elements. Select all of the algorithms that you want to allow in the VPN. For a full explanation of the options, see Deciding Which Hash Algorithms to Use, on page 1503.

Step 10 Choose a **Diffie-Hellman Group** for encryption.

A larger modulus provides higher security but requires more processing time. The two peers must have a matching modulus group. Select the groups that you want to allow in the VPN. For a full explanation of the options, see Deciding Which Diffie-Hellman Modulus Group to Use, on page 1503.

Step 11 Click **Save**
If a valid combination of choices has been selected the new IKEv2 policy is added to the list. If not, errors are displayed and you must make changes accordingly to successfully save this policy.

Secure Client Customization

You can now customize Secure Client and deploy these customizations to the VPN headend. The threat defense distributes these customizations to the endpoint when an end user connects from the Secure Client.

Secure Client Customization objects represent files used to customize Secure Client. The following are the supported Secure Client customizations:

- GUI text and messages
- Icons and images
- Scripts
- Binaries
- Custom Installer Transforms
- Localized Installer Transforms

For more information about configuring these Secure Client customizations, see Customize Cisco Secure Client, on page 1629.

File Objects

Use the Add and Edit File Object dialog boxes to create, and edit file objects. File objects represent files used in configurations, typically for remote access VPN policies. They can contain Secure Client Profile and Secure Client Image files.

Profiles are also created for each AnyConnect module and Secure Client Management VPN using independent profile editors and deployed to administrator-defined end user requirements and authentication policies on endpoints as part of Secure Client, and they make the preconfigured network profiles available to end users.

When you create a file object, the Firewall Management Center makes a copy of the file in its repository. These files are backed up whenever you create a backup of the database, and they are restored if you restore the database. When copying a file to the management center platform to be used in a file object, do not copy the file directly to the file repository.

When you deploy configurations that specify a file object, the associated file is downloaded to the device in the appropriate directory.

You can click one of the following options against each file:

- **Download** —Click to download the Secure Client file.
- **Edit** —Modify the file object details.
- **Delete** —Delete the Secure Client file object. When you delete a file object, the associated file is not deleted from the file repository, only the object is deleted.

Navigation Path

Objects > **Object Management** > **VPN** > **Secure Client File**.

Fields

- **Name**—Enter the name of the file to identify the file object; you can add up to 128 characters.

- **File Name**—Click **Browse** to select the file. The file name and full path of the file are added when you select the file.
- **File Type**—Choose the file type corresponding to the file you have selected. The following file types are available:
  - **Secure Client Image**—Select this type when you add the Secure Client image you have downloaded from the Cisco Software Download Center.

    You can associate any new or additional Secure Client images to the remote access VPN policy. You can also unassociate the unsupported or end of life client packages that are no longer required.
  - **Secure Client VPN Profile**—Choose this type for the Secure Client VPN profile file.

    The profile file is created using the GUI-based Secure Client Profile Editor, an independent configuration tool. See the *Secure Client Profile Editor* chapter in the appropriate release of the Cisco Secure Client (including AnyConnect) User Guide for details.
  - **Secure Client Management VPN Profile**—Select this type when you add a profile file for the Secure Client management VPN tunnel.

    Download the Secure Client **VPN Management Tunnel Standalone Profile Editor** from Cisco Software Download Center if you have not done already and create a profile with required settings for the Secure Client management VPN tunnel.
  - **AMP Enabler Service Profile**—The profile is used for the Secure Client AMP Enabler. The AMP Enabler along with this profile is pushed to the endpoints from Firewall Threat Defense when a remote access VPN user connects to the VPN.
  - **Feedback Profile**—You can add a Customer Experience Feedback profile and select this type to receive information about the features and modules customers have enabled and use.
  - **ISE Posture Profile**—Choose this option if you are adding a profile file for the Secure Client ISE Posture module.
  - **NAM Service Profile**—Configure and add the NAM profile file using the Network Access Manager profile editor.
  - **Network Visibility Service Profile**—Profile file for Secure Client Network Visibility module. You can create the profile using the NVM profile editor.
  - **Umbrella Roaming Security Profile**—You must select this file type if you are deploying the Umbrella Roaming Security module using the .json file created using the profile editor.
  - **Web Security Service Profile**—Select this file type when you add a prole file for the Web security module.
  - **Secure Firewall Posture Package**—Select this file type when you add a Secure Firewall Posture Package file. This file is used while configuring a Dynamic Access Policy (DAP) to collect information about the operating system, anti-virus, anti-spyware, and firewall software installed on the endpoints.
  - **Secure Client External Browser Package**—This file type is for selecting an external browser package file for SAML single sing-on web authentication.

    You can add an the package file when a new version of the external package file is available.

    For more information, see Configure AAA Settings for Remote Access VPN, on page 1597.

- **Description**—Add an optional description.

Related Topics

Cisco Secure Client Image, on page 1615
Group Policy Secure Client Options, on page 1469

History for Object Management

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Merged ACL and AV ACL	7.4.1	Any	New/modified screens: **Objects** > > **Object Management** > **RADIUS Server Group** > **Add RADIUS Server Group** > **Merge Downloadable ACL with Cisco AV Pair ACL** New CLI commands: • **sh run aaa-server aaa-server ISE-Server protocol radius merge-dacl after-avpair** • **sh run aaa-server aaa-server ISE-Server protocol radius merge-dacl before-avpair**
Secure Client Customization	Any	7.4	You can now customize Secure Client and deploy these customizations to the VPN headend. The threat defense distributes these customizations to the endpoint when an end user connects from the Secure Client.
IPv6 support for CRL and OCSP urls	Any	7.4	You can now configure IPv6 OCSP and CRL URLs.
Loopback and Management type interface group objects	Any	7.4	You can now create interface group objects that include only management-only interfaces or only loopback interfaces. You can then use these groups for management features such as DNS servers, HTTP access, or SSH. Loopback groups are supported for any feature that supports loopback intefaces. Note that DNS does not support management interfaces. New/Modified screens: **Objects** > **Object Management** > **Interface** > **Add** > **Interface Group**
Loopback interface support for AAA	Any	7.4	You can use a loopback interface group for configuring a RADIUS server. New/Modified screens: **Objects** > **Object Management** > **AAA Server** > **RADIUS Server Group**
Clone network and port objects	Any	7.4	You can now clone network and port objects. In the object manager (**Objects** > **Object Management**), click the new **Clone** icon next to a port or network object. You can then change the new object's properties and save it using a new name.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
DHCP IPv6 Pool	Any	7.3	The Firewall Threat Defense now supports a light DHCPv6 stateless server when using the DHCPv6 Prefix Delegation client. The Firewall Threat Defense provides other information such as the domain name to SLAAC clients when they send Information Request (IR) packets to the Firewall Threat Defense. The Firewall Threat Defense only accepts IR packets and does not assign addresses to the clients. New/Modified screens: • **Devices** > **Device Management** > **Interfaces** > **Add/Edit Interfaces** > **IPv6** > **DHCP** • **Objects** > **Object Management** > **DHCP IPv6 Pool** New/Modified commands: **show ipv6 dhcp**
BFD Template	Any	7.3	In the previous releases, BFD was configurable on threat defense only through FlexConfig. FlexConfig no longer supports BFD configuration. You can now configure BFD policies for threat defense in the management center UI. Hence, the BFD Template object was introduced. New/modified screens: • **Devices** > **Device Management** > **Routing** > **BFD**. • **Objects** > **Object Management** > **BFD Template**
New Applications tab for policy based routing	Any	7.1	A new tab for selecting the applications for configuring direct internet access policy (policy based routing) was introduced in the extended access list object. New/Modified Screens: New option to select applications when configuring **Object** > **Object Management** > **Access List** > **Extended** page. Supported platforms: Secure Firewall Management Center
New Extended Community List object and	Any	7.1	Extended community list object was introduced to be used in policy list and route map objects. The extended community list object is applicable only for importing or exporting routes in BGP route leaking support for virtual routers. New/Modified Screens: New object for configuring policy list and route maps **Object** > **Object Management** > **Community list** > **Extended Community** page. Supported platforms: Secure Firewall Management Center

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Enhancements to Policy List object and Route Map object	Any	7.1	Options to select the newly introduced Extended Community List objects in policy list and route maps. New/Modified Screens: New options for configuring policy list and route maps **Object** > **Object Management** > **Policy List** > **Community rule** tab, and **Object** > **Object Management** > **Route Map** > **BGP** > **Community List** tab. Supported platforms: Secure Firewall Management Center
Time-based ACL support for Snort 3	Any	7.0	Time-based rules in access control and prefilter policies are supported in Snort 3 as well. Supported platforms: Firewall Threat Defense
EST for certificate enrollment	Any	7.0	Support for Enrollment over Secure Transport for certificate enrollment was provided. New/Modified Screens: New enrollment options when configuring **Objects** > **PKI** > **Cert Enrollmet** > **CA Information** tab. Supported platforms: Secure Firewall Management Center
Support for EdDSA certificate type	Any	7.0	A new certificate key type- EdDSA was added with key size 256. New/Modified Screens: New certificate key options when configuring **Objects** > **PKI** > **Cert Enrollmet** > **Key** tab. Supported platforms: Secure Firewall Management Center
Restrictions on ciphers and key sizes	Any	7.0	Certificates having SHA-1 with RSA Encryption signature algorithm, and RSA key sizes smaller than 2048 bits are not supported. To override these restrictions on existing certificates, you can enable the weak-crypto option on Firewall Threat Defense. However, you cannot generate RSA keys with sizes smaller than 2048 bits. New/Modified Screens: New toggle button when configuring **Devices** > **Certificates**. Supported platforms: Secure Firewall Management Center
Security Intelligence feed options	Any	6.7	New update frequency options (5 and 15 minutes) for custom Security Intelligence feeds. Update frequencies of less than 30 minutes require an MD5 URL, to prevent unnecessary downloads if the feed has not changed. New/Modified Screens: New frequency choices when configuring **Security Intelligence** > **Network Lists and Feeds**. Supported platforms: Secure Firewall Management Center

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Bulk upload of objects using a comma-separated-values (csv) file	Any	6.7	Objects can be imported from a comma-separated-values file. Up to 1000 objects can be imported in one attempt. New/Modified Screens: The following object types have a new **Import Object** option in the **Add [Object Type]** drop-down list. • **Distinguished Name** > **Individual Objects** • **Network Object** • **Port** • **URL** • **VLAN Tag** Supported platforms: Secure Firewall Management Center
See the policies in which interface objects are used	Any	6.6	See the policies in which interface objects are used. New/Modified Screens: The **Interface** object page in **Object** > **Object management** has a new **Find** button. Supported platforms: Secure Firewall Management Center
Time zone objects introduced	Any	6.6	You can assign time zones to Firewall Threat Defense devices, for use when applying time-based policies. New/Modified Screens: New **Time Zone Object** in **Object** > **Object management**. Supported platforms: Secure Firewall Management Center
Time-based objects can now be used in access control and prefilter policies	Any	6.6	Use time range objects in conjunction with new time zone objects for applying time-based rules in access control and prefilter policies. You can specify an absolute or recurring time or time range for a rule to be applied. The rule is applied based on the time zone of the device that processes the traffic.
View Object Details from prefilter rule page	Any	6.6	Feature introduced: Option to view details for an object or object group when viewing prefilter rules. New options: Right-clicking a value in any of the following columns in the prefilter rule list offers options to view object details: Source Networks, Destination Networks, Source Port, Destination Port, and VLAN Tag. Supported platforms: Secure Firewall Management Center

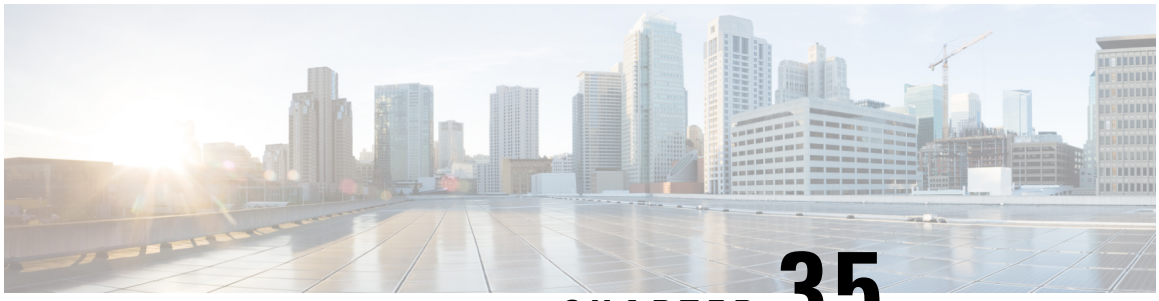

CHAPTER 35

Certificates

- Requirements and Prerequisites for Certificates, on page 1485
- Secure Firewall Threat Defense VPN Certificate Guidelines and Limitations, on page 1485
- Managing Firewall Threat Defense Certificates, on page 1486
- Installing a Certificate Using Self-Signed Enrollment , on page 1489
- Installing a Certificate using EST Enrollment, on page 1490
- Installing a Certificate Using SCEP Enrollment, on page 1491
- Installing a Certificate Using Manual Enrollment, on page 1491
- Installing a Certificate Using a PKCS12 File, on page 1492
- Troubleshooting Firewall Threat Defense Certificates, on page 1493
- History for Certificates, on page 1494

Requirements and Prerequisites for Certificates

Supported Domains

Any

User Roles

Admin

Network Admin

Secure Firewall Threat Defense VPN Certificate Guidelines and Limitations

- When a PKI enrollment object is associated with and then installed on a device, the certificate enrollment process starts immediately. The process is automatic for self-signed and SCEP enrollment types; it does not require any additional administrator's action. Manual certificate enrollment requires administrator's action.
- When the certificate enrollment is complete, a trustpoint exists on the device with the same name as the certificate enrollment object. Use this trustpoint in the configuration of your VPN Authentication Method.

- Firewall Threat Defense devices support certificate enrollment using Microsoft Certificate Authority(CA) Service, and CA Services provided on Cisco Adaptive Security Appliances(ASA) and Cisco IOS Router.
- Firewall Threat Defense devices cannot be configured as a certificate authority (CA).

Guidelines for Certificate Management Across Domains and Devices

- Certificate enrollment can be done in a child or parent domain.
- When enrollment is done from a parent domain, the certificate enrollment object also needs to be in the same domain. If the trustpoint on a device is overridden in the child domain, the overridden value will be deployed on the device.
- When the certificate enrollment is done on a device in a leaf domain, the enrollment will be visible to the parent domain or another child domain. Also, adding additional certificates is possible.
- When a leaf domain is deleted, certificate enrollments on the contained devices will be automatically removed.
- Once a device has certificates enrolled in one domain, it will be allowed to be enrolled in any other domain. The certificates can be added in the other domain.
- When you move a device from one domain to another, the certificates also get moved accordingly. You will receive an alert to delete the enrollments on these devices.

Managing Firewall Threat Defense Certificates

See PKI Infrastructure and Digital Certificates , on page 1505 for an introduction to Digital Certificates.

See Certificate Enrollment Objects, on page 1413 for a description of the objects used to enroll and obtain certificates on managed devices.

Procedure

Step 1 Select **Devices** > **Certificates**.

You can see the following columns for each device listed on this screen:

- **Name**—Lists the devices that already have trustpoints associated with them. Expand the device to see the list of associated trustpoints.
- **Domain**—Displays the certificates that are enrolled in a specific domain.
- **Enrollment Type**—Displays the type of enrollment used for a trustpoint.
- **Status**—Provides the status of the **CA Certificate** and **Identity Certificate**. You can view the certificate contents, when `Available`, by clicking the magnifying glass.

  When you view the CA certificate information, you can view the hierarchy of all the certifying authorities, which issued your CA certificate. If the enrollment fails, click status to view the failure message.
- Click **Enable weak-crypto** on the right to enable weak cipher usage in certificates. When you click the toggle button, you get a warning to confirm before enabling weak ciphers. Click **Yes** to enable weak ciphers.

Note
When a certificate enrollment fails due to weak cipher usage, you get a prompt to enable the weak cipher. You can choose to enable weak cipher when you need to use weak encryption.

- The additional column lists icons to perform the following tasks:
  - **Export Certificate**—Click to export and download a copy of the certificate. You can choose to export the PKCS12 (Complete Certificate Chain) or the PEM (Identity Certificate Only) format.

    You must provide a pass phrase to export a PKCS12 certificate format to import the file later.
  - **Re-enroll certificate**—Re-enroll an existing certificate.
  - **Refresh certificate status**—Refresh a certificate to synchronize the device certificate status to the Firewall Management Center.
  - **Delete certificate**—Delete all the associated certificates for a trustpoint.

Step 2 Choose **(+) Add** to associate and install an enrollment object on a device.

When a certificate enrollment object is associated with and then installed on a device, the process of certificate enrollment starts immediately. The process is automatic for self-signed and SCEP enrollment types, meaning it does not require any additional administrator action. Manual certificate enrollment requires extra administrator action.

Note
The certificate enrollment on a device does not block the user interface and the enrollment process gets executed in the background, enabling the user to perform certificate enrollment on other devices in parallel. The progress of these parallel operations can be monitored on the same user interface. The respective icons display the certificate enrollment status.

Related Topics

Installing a Certificate Using Self-Signed Enrollment , on page 1489
Installing a Certificate Using SCEP Enrollment, on page 1491
Installing a Certificate Using Manual Enrollment, on page 1491
Installing a Certificate Using a PKCS12 File, on page 1492

Automatically Update CA Bundles

You can set the management center to automatically update the CA certificates through CLI commands. By default, the CA certificates are automatically updated when you install or upgrade to version 7.0.5.

Note In an IPv6-only deployment, the automatic update of CA certificates may fail, because, some of the Cisco servers do not support IPv6. In such cases, force update the CA certificates using the **configure cert-update run-now force** command.

Procedure

Step 1 Log into the FMC CLI using SSH, or, if virtual, open the VM console.

Step 2 You can verify whether the CA certificates in the local system are the latest or not:

configure cert-update test

This command compares the CA bundle on the local system with the latest CA bundle (from the Cisco server). If the CA bundle is up to date, no connection check is executed and the test result is displayed as the one below:

Example:

```
> configure cert-update test
Test succeeded, certs can safely be updated or are already up to date.
```

If the CA bundle is out of date, the connection check is executed on the downloaded CA bundle and the test result is displayed.

Example:

When the connection check fails:

```
> configure cert-update test
Test failed, not able to fully connect.
```

Example:

When the connection check succeeds, or the CA bundle is already up to date:

```
> configure cert-update test
Test succeeded, certs can safely be updated or are already up to date.
```

Step 3 (Optional) To instantly update the CA bundles:

configure cert-update run-now

Example:

```
>configure cert-update run-now
Certs have been replaced or was already up to date.
```

When you execute this command, the CA certificates (from the Cisco server) are verified for SSL connectivity. If the SSL connectivity check fails for even one of the Cisco servers, the process is terminated.

Example:

```
> configure cert-update run-now
Certs failed some connection checks.
```

To proceed with the update despite connection failures, use the **force** keyword.

Example:

```
> configure cert-update run-now force
```

```
Certs failed some connection checks, but replace has been forced.
```

Step 4 If you do not want the CA bundles to be automatically updated, disable the configuration:

configure cert-update auto-update disable

Example:

```
> configure cert-update auto-update disable
Autoupdate is disabled
```

Step 5 To re-enable the automatic update of CA bundles:

configure cert-update auto-update enable

Example:

```
> configure cert-update auto-update enable
Autoupdate is enabled and set for every day at 12:18 UTC
```

When you enable the automatic update on the CA certificates, the update process is executed daily at a system-defined time.

Step 6 (Optional) View the status of automatic update of CA certificates:

show cert-update

Example:

```
> show cert-update
Autoupdate is enabled and set for every day at 09:34 UTC
CA bundle was last modified 'Thu Sep 15 16:12:35 2022'
```

Installing a Certificate Using Self-Signed Enrollment

Procedure

Step 1 On the **Devices** > **Certificates** screen, choose **Add** to open the **Add New Certificate** dialog.

Step 2 Choose a device from the **Device** drop-down list.

Step 3 Associate a certificate enrollment object with this device in one of the following ways:

- Choose a Certificate Enrollment Object of the type Self-Signed from the drop-down list.
- Click **(+)**, to add a new Certificate Enrollment Object, see Adding Certificate Enrollment Objects, on page 1414.

Step 4 Press **Add** to start the Self Signed, automatic, enrollment process.

For self signed enrollment type trustpoints, the **CA Certificate** status will always be displayed, since the managed device is acting as its own CA and does not need a CA certificate to generate its own Identity Certificate.

The **Identity Certificate** will go from InProgress to Available as the device creates its own self signed identity certificate.

Step 5 Click the magnifying glass to view the self-signed Identity Certificate created for this device.

What to do next

When enrollment is complete, a trustpoint exists on the device with the same name as the certificate enrollment object. Use this trustpoint in the configuration of your Site to Site and Remote Access VPN Authentication Method

Installing a Certificate using EST Enrollment

Before you begin

Note Using EST enrollment establishes a direct connection between the managed device and the CA server. So be sure your device is connected to the CA server before beginning the enrollment process.

Note EST's ability to auto-enroll a device when its certificate expires is not supported.

Procedure

Step 1 On the **Devices** > **Certificates** screen, click **Add** to open the **Add New Certificate** dialog.

Step 2 Choose a device from the **Device** drop-down list.

Step 3 Associate a certificate enrollment object with this device in one of the following ways:

- Choose the EST certificate enrollment object from the **Cert Enrollment** drop-down list.
- Click **(+)**, to add a new Certificate Enrollment Object, see Adding Certificate Enrollment Objects, on page 1414.

Step 4 Click **Add** to enroll the certificate on the device.

The **Identity Certificate** will go from `InProgress` to `Available` as the device obtains its identity certificate using EST from the specified CA. Sometimes, a manual refresh might be required to obtain the identity certificate.

Step 5 Click the magnifying glass to view the Identity Certificate created and installed on this device.

Installing a Certificate Using SCEP Enrollment

Before you begin

Note Using SCEP enrollment establishes a direct connection between the managed device and the CA server. So be sure your device is connected to the CA server before beginning the enrollment process.

Procedure

Step 1 On the **Devices** > **Certificates** screen, choose **Add** to open the **Add New Certificate** dialog.

Step 2 Choose a device from the **Device** drop-down list.

Step 3 Associate a certificate enrollment object with this device in one of the following ways:

- Choose a Certificate Enrollment Object of the type SCEP from the drop-down list.
- Click **(+)**, to add a new Certificate Enrollment Object, see Adding Certificate Enrollment Objects, on page 1414.

Step 4 Press **Add**, to start the automatic enrollment process.

For SCEP enrollment type trustpoints, the **CA Certificate** status will transition from InProgress to Available as the CA Certificate is obtained from the CA server and installed on the device.

The **Identity Certificate** will go from `InProgress` to `Available` as the device obtains its identity certificate using SCEP from the specified CA. Sometimes, a manual refresh might be required to obtain the identity certificate.

Step 5 Click the magnifying glass to view the Identity Certificate created and installed on this device.

What to do next

When enrollment is complete, a trustpoint exists on the device with the same name as the certificate enrollment object. Use this trustpoint in the configuration of your Site to Site and Remote Access VPN Authentication Method

Installing a Certificate Using Manual Enrollment

Procedure

Step 1 On the **Devices** > **Certificates** screen, choose **Add** to open the **Add New Certificate** dialog.

Step 2 Choose a device from the **Device** drop-down list.

Step 3 Associate a certificate enrollment object with this device in one of the following ways:

- Choose a Certificate Enrollment Object of the type Manual from the drop-down list.
- Click **(+)**, to add a new Certificate Enrollment Object, see Adding Certificate Enrollment Objects, on page 1414.

Step 4 Press **Add** to start the enrollment process.

Step 5 Execute the appropriate activity with your PKI CA Server to obtain an identity certificate.

a) Click **Identity Certificate** warning to view and copy the CSR.

b) Execute the appropriate activity with your PKI CA Server to obtain an identity certificate using this CSR.

This activity is completely independent of the Secure Firewall Management Center or the managed device. When complete, you will have an Identity Certificate for the managed device. You can place it in a file.

c) To finish the manual process, install the obtained identity certificate onto the managed device.

Return to the Secure Firewall Management Center dialog and select **Browse Identity Certificate** to choose the identity certificate file.

Note
Ensure not to choose a binary certificate (PKCS12, DER, and alike) file because Firewall Threat Defense does not support them.

Step 6 Select **Import** to import the Identity Certificate.

The Identity Certificate status will be `Available` when the import complete.

Step 7 Click the magnifying glass to view the **Identity Certificate** for this device.

What to do next

When enrollment is complete, a trustpoint exists on the device with the same name as the certificate enrollment object. Use this trustpoint in the configuration of your Site to Site and Remote Access VPN Authentication Method

Installing a Certificate Using a PKCS12 File

Procedure

Step 1 Go to **Devices** > **Certificates** screen, choose **Add** to open the **Add New Certificate** dialog.

Step 2 Choose a pre-configured managed device from the **Device** drop down list.

Step 3 Associate a certificate enrollment object with this device in one of the following ways:

- Choose a Certificate Enrollment Object of the PKCS type from the drop-down list.
- Click **(+)**, to add a new Certificate Enrollment Object, see Adding Certificate Enrollment Objects, on page 1414.

Step 4 Press **Add**.

The CA Certificate and Identity Certificate status will go from `In Progress` to `Available` as it installs the PKCS12 file on the device.

Note
When you upload the PKCS12 file for the first time, the file is stored in Firewall Management Center as part of the CertEnrollment object. For any failed enrollments due to a wrong passphrase or failed deployment, retry enrolling the PKCS12 certificate without uploading the file again. Also, whenever you modify the certificate enrollment object to permit overrides, you must update the **Passphrase** for the certificate for a sucessful enrollment. A PKCS12 file size should not be larger than 24K.

Step 5 Once `Available`, click the magnifying glass to view the Identity Certificate for this device.

What to do next

The certificate (trustpoint) on the managed device is named the same as the PKCS#12 file. Use this certificate in your VPN authentication configuration.

Troubleshooting Firewall Threat Defense Certificates

See Secure Firewall Threat Defense VPN Certificate Guidelines and Limitations, on page 1485 to determine if variations in your certificate enrollment environment may be causing a problem. Then consider the following:

- **Symptom**: Expiration of service authentication certificates.

  A certificate monitoring health module alerts you before service authentication certificates expire on the Firewall Management Center and managed devices. Choose **System** (⚙) > **Health** > **Policy**, click the **Health Modules** tab and then navigate to the **Certificate Monitoring** tile.

- Ensure there is a route to the CA Server from the device.

  If the CA Server's host name is given in the Enrollment Object, use Flex Config to configure DNS appropriately to reach the server. Alternatively, use the IP Address of the CA Server.

- If you are using a Microsoft 2012 CA Server, the default IPsec Template is not accepted by the managed device and must be changed.

  To configure a working template, follow these steps as you use MS CA documentation as a reference.

  1. Duplicate the IPsec (Offline Request) template.
  2. In **Extensions > Application policies**, select *IP security end system*, instead of the *IP security IKE intermediate*.
  3. Set the permissions and the template name.
  4. Add the new template and change the registry settings to reflect the new template name.

- On the Firewall Management Center, you might receive the following health alert related to the Firewall Threat Defense device:

  ```
  Code - F0853; Description - default Keyring's certificate is invalid, reason: expired
  ```

  Solution: In such cases, use the following command to regenerate the default certificate in CLISH CLI:

  ```
  > system support regenerate-security-keyring default
  ```

- A red cross appears in the CA certificate status with the following error:

  ```
  Fail to configure CA certificate
  ```

Solution: See Troubleshoot Certificate Error on FMC.

- To check the list of the certificates in a .pfx file, use tools such as `certutil` or `openssl`. You can see the whole chain with ID certificate, SubCA certificate, and CA certificate (if any).
  - `certutil -dump cert.pfx`
  - `openssl pkcs12 -info -in cert.pfx`
- The following error appears:

  `Identity certificate import required`

  Solution: See Troubleshoot Certificate Error "Identity certificate import required" on FMC.

History for Certificates

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Support for OCSP and CRL IPv6 URL	7.4	Any	You can now add IPv6 OCSP and CRL URLs for certificate authentication (revocation checking). The IPv6 addresses must be enclosed in square brackets.
Enhancements to Manual Enrollment	6.7	Any	You can now create only a CA certificate without an identity certificate. You can also generate a CSR without a CA certificate and obtain an identity certificate from the CA.
PKCS CA Chain	6.7	Any	You can view and manage the chain of certifying authorities (CAs) issuing your certificates. You can also export a copy of the certificates.

PART VI

VPN

- VPN Overview, on page 1497
- Site-to-Site VPNs, on page 1511
- Remote Access VPN, on page 1573
- Dynamic Access Policies , on page 1689
- VPN Monitoring and Troubleshooting, on page 1703

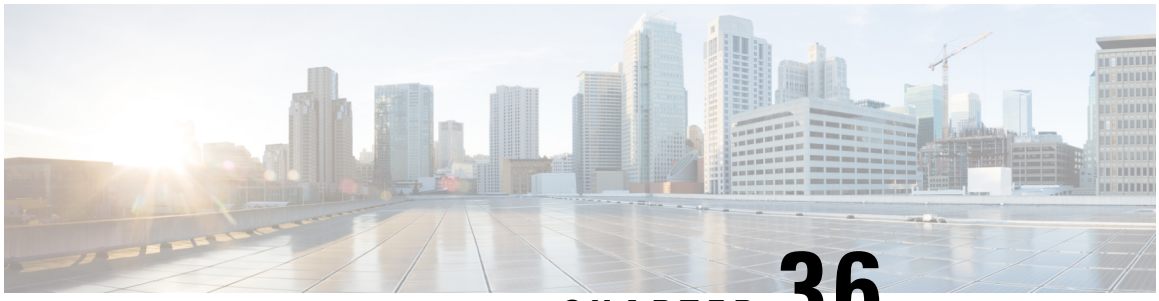

CHAPTER 36

VPN Overview

A virtual private network (VPN) connection establishes a secure tunnel between endpoints over a public network such as the Internet.

This chapter applies to Remote Access and Site-to-site VPNs on Secure Firewall Threat Defense devices. It describes the Internet Protocol Security (IPsec), the Internet Security Association and Key Management Protocol (ISAKMP, or IKE) and SSL standards that are used to build site-to-site and remote access VPNs.

- VPN Types, on page 1497
- VPN Basics, on page 1498
- VPN Packet Flow, on page 1500
- IPsec Flow Offload, on page 1500
- VPN Licensing, on page 1501
- How Secure Should a VPN Connection Be?, on page 1502
- Removed or Deprecated Hash Algorithms, Encryption Algorithms, and Diffie-Hellman Modulus Groups, on page 1506
- VPN Topology Options, on page 1507

VPN Types

The Firewall Management Center supports the following types of VPN connections:

- Remote Access VPNs on Firewall Threat Defense devices.

  Remote access VPNs are secure, encrypted connections, or tunnels, between remote users and your company's private network. The connection consists of a VPN endpoint device, which is a workstation or mobile device with VPN client capabilities, and a VPN headend device, or secure gateway, at the edge of the corporate private network.

  Secure Firewall Threat Defense devices can be configured to support Remote Access VPNs over SSL or IPsec IKEv2 by the Firewall Management Center. Functioning as secure gateways in this capacity, they authenticate remote users, authorize access, and encrypt data to provide secure connections to your network. No other types of appliances, managed by the Firewall Management Center, support Remote Access VPN connections.

  Secure Firewall Threat Defense secure gateways support the Secure Client full tunnel client. This client is required to provide secure SSL IPsec IKEv2 connections for remote users. This client gives remote users the benefits of a client without the need for network administrators to install and configure clients

on remote computers since it can be deployed to the client platform upon connectivity. It is the only client supported on endpoint devices.

- Site-to-site VPNs on Firewall Threat Defense devices.

  A site-to-site VPN connects networks in different geographic locations. You can create site-to-site IPsec connections between managed devices, and between managed devices and other Cisco or third-party peers that comply with all relevant standards. These peers can have any mix of inside and outside IPv4 and IPv6 addresses. Site-to-site tunnels are built using the Internet Protocol Security (IPsec) protocol suite and IKEv1 or IKEv2. After the VPN connection is established, the hosts behind the local gateway can connect to the hosts behind the remote gateway through the secure VPN tunnel.

VPN Basics

Tunneling makes it possible to use a public TCP/IP network, such as the Internet, to create secure connections between remote users and private corporate networks. Each secure connection is called a tunnel.

IPsec-based VPN technologies use the Internet Security Association and Key Management Protocol (ISAKMP, or IKE) and IPsec tunneling standards to build and manage tunnels. ISAKMP and IPsec accomplish the following:

- Negotiate tunnel parameters.
- Establish tunnels.
- Authenticate users and data.
- Manage security keys.
- Encrypt and decrypt data.
- Manage data transfer across the tunnel.
- Manage data transfer inbound and outbound as a tunnel endpoint or router.

A device in a VPN functions as a bidirectional tunnel endpoint. It can receive plain packets from the private network, encapsulate them, create a tunnel, and send them to the other end of the tunnel where they are unencapsulated and sent to their final destination. It can also receive encapsulated packets from the public network, unencapsulate them, and send them to their final destination on the private network.

After the site-to-site VPN connection is established, the hosts behind the local gateway can connect to the hosts behind the remote gateway through the secure VPN tunnel. A connection consists of the IP addresses and hostnames of the two gateways, the subnets behind them, and the method the two gateways use to authenticate to each other.

Hubs—Devices that enable secure VPN connectivity to and from one or more remote branch devices or spokes. Hubs also act as a gateway for spokes to communicate with each other.

Spokes—Devices that connect over VPN to a hub to securely access the corporate resources behind the hub. Spokes communicate with each other through the hub.

Internet Key Exchange (IKE)

Internet Key Exchange (IKE) is a key management protocol that is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and to automatically establish IPsec security associations (SAs).

The IKE negotiation comprises two phases. Phase 1 negotiates a security association between two IKE peers, which enables the peers to communicate securely in Phase 2. During Phase 2 negotiation, IKE establishes SAs for other applications, such as IPsec. Both phases use proposals when they negotiate a connection.

An IKE policy is a set of algorithms that two peers use to secure the IKE negotiation between them. IKE negotiation begins by each peer agreeing on a common (shared) IKE policy. This policy states which security parameters protect subsequent IKE negotiations. For IKE version 1 (IKEv1), IKE policies contain a single set of algorithms and a modulus group. Unlike IKEv1, in an IKEv2 policy, you can select multiple algorithms and modulus groups from which peers can choose during the Phase 1 negotiation. It is possible to create a single IKE policy, although you might want different policies to give higher priority to your most desired options. For site-to-site VPNs, you can create an IKE policy. IKEv1 and IKEv2 each support a maximum of 20 IKE policies, each with a different set of values. Assign a unique priority to each policy that you create. The lower the priority number, the higher the priority.

To define an IKE policy, specify:

- A unique priority (1 to 65,543, with 1 the highest priority).
- An encryption method for the IKE negotiation, to protect the data and ensure privacy.
- A Hashed Message Authentication Codes (HMAC) method (called integrity algorithm in IKEv2) to ensure the identity of the sender, and to ensure that the message has not been modified in transit.
- For IKEv2, a separate pseudorandom function (PRF) used as the algorithm to derive keying material and hashing operations required for the IKEv2 tunnel encryption. The options are the same as those used for the hash algorithm.
- A Diffie-Hellman group to determine the strength of the encryption-key-determination algorithm. The device uses this algorithm to derive the encryption and hash keys.
- An authentication method, to ensure the identity of the peers.
- A limit to the time the device uses an encryption key before replacing it.

When IKE negotiation begins, the peer that starts the negotiation sends all of its policies to the remote peer, and the remote peer searches for a match with its own policies, in priority order. A match between IKE policies exists if they have the same encryption, hash (integrity and PRF for IKEv2), authentication, and Diffie-Hellman values, and an SA lifetime less than or equal to the lifetime in the policy sent. If the lifetimes are not identical, the shorter lifetime—From the remote peer policy—Applies. By default, the Secure Firewall Management Center deploys an IKEv1 policy at the lowest priority for all VPN endpoints to ensure a successful negotiation.

IPsec

IPsec is one of the most secure methods for setting up a VPN. IPsec provides data encryption at the IP packet level, offering a robust security solution that is standards-based. With IPsec, data is transmitted over a public network through tunnels. A tunnel is a secure, logical communication path between two peers. Traffic that enters an IPsec tunnel is secured by a combination of security protocols and algorithms.

An IPsec Proposal policy defines the settings required for IPsec tunnels. An IPsec proposal is a collection of one or more crypto-maps that are applied to the VPN interfaces on the devices. A crypto map combines all the components required to set up IPsec security associations, including:

- A proposal (or transform set) is a combination of security protocols and algorithms that secure traffic in an IPsec tunnel. During the IPsec security association (SA) negotiation, peers search for a proposal that is the same at both peers. When it is found, it is applied to create an SA that protects data flows in the access list for that crypto map, protecting the traffic in the VPN. There are separate IPsec proposals for IKEv1 and IKEv2. In IKEv1 proposals (or transform sets), for each parameter, you set one value. For IKEv2 proposals, you can configure multiple encryption and integration algorithms for a single proposal.
- A crypto map, combines all components required to set up IPsec security associations (SA), including IPsec rules, proposals, remote peers, and other parameters that are necessary to define an IPsec SA. When two peers try to establish an SA, they must each have at least one compatible crypto map entry.

  Dynamic crypto map policies are used in site-to-site VPNs when an unknown remote peer tries to start an IPsec security association with the local hub. The hub cannot be the initiator of the security association negotiation. Dynamic crypto-policies allow remote peers to exchange IPsec traffic with a local hub even if the hub does not know the remote peer's identity. A dynamic crypto map policy essentially creates a crypto map entry without all the parameters configured. The missing parameters are later dynamically configured (as the result of an IPsec negotiation) to match a remote peer's requirements.

  Dynamic crypto map policies are applicable to both hub-and-spoke and point-to-point VPN topologies. To apply dynamic crypto map policies, specify a dynamic IP address for one of the peers in the topology and ensure that the dynamic crypto-map is enabled on this topology. Note that in a full mesh VPN topology, you can apply only static crypto map policies.

Note Simultaneous IKEv2 dynamic crypto map is not supported for the same interface for both remote access and site-to-site VPNs on Firewall Threat Defense.

VPN Packet Flow

On a Firewall Threat Defense device, by default no traffic is allowed to pass through access-control without explicit permission. VPN tunnel traffic as well, is not relayed to the endpoints until it has passed through Snort. Incoming tunnel packets are decrypted before being sent to the Snort process. Snort processes outgoing packets before encryption.

Access Control identifying the protected networks for each endpoint node of a VPN tunnel determines which traffic is allowed to pass through the Firewall Threat Defense device and reach the endpoints. For Remote Access VPN traffic, a Group Policy filter or an Access Control rule must be configured to permit VPN traffic flow.

In addition, the system does not send tunnel traffic to the public source when the tunnel is down.

IPsec Flow Offload

You can configure supporting device models to use IPsec flow offload. After the initial setup of an IPsec site-to-site VPN or remote access VPN security association (SA), IPsec connections are offloaded to the

field-programmable gate array (FPGA) in the device, which should improve device performance. On the Secure Firewall 1200 series, IPsec connections are offloaded to the Marvell Cryptographic Accelerator (CPT) to improve device performance.

Offloaded operations specifically relate to the pre-decryption and decryption processing on ingress, and the pre-encryption and encryption processing on egress. The system software handles the inner flow to apply your security policies.

IPsec flow offload is enabled by default, and applies to the following device types:

- Secure Firewall 1200
- Secure Firewall 3100
- Secure Firewall 4200

IPsec flow offload is also used when the device's VTI loopback interface is enabled.

Limitations for IPsec Flow Offload

The following IPsec flows are not offloaded:

- IKEv1 tunnels. Only IKEv2 tunnels will be offloaded. IKEv2 supports stronger ciphers.
- Flows that have volume-based rekeying configured.
- Flows that have compression configured.
- Transport mode flows. Only tunnel mode flows will be offloaded.
- AH format. Only ESP/NAT-T format will be supported.
- Flows that have post-fragmentation configured.
- Flows that have anti-replay window size other than 64bit and anti-replay is not disabled.
- Flows that have firewall filter enabled.
- Mult-instance mode.

Configure IPsec Flow Offload

IPsec flow offload is enabled by default on hardware platforms that support the feature. To change the configuration, use FlexConfig to implement the **flow-offload-ipsec** command. See the ASA command reference for detailed information about the command.

VPN Licensing

There is no specific licensing for enabling Secure Firewall Threat Defense VPN, it is available by default.

The Firewall Management Center determines whether to allow or block the usage of strong crypto on the Firewall Threat Defense device based on attributes provided by the smart licensing server.

This is controlled by whether you selected the option to allow export-controlled functionality on the device when you registered with the Cisco Smart License Manager. If you are using the evaluation license, or you did not enable export-controlled functionality, you cannot use strong encryption.

If you have created your VPN configurations with an evaluation license, and upgrade your license from evaluation to smart license with export-controlled functionality, check, and update your encryption algorithms for stronger encryption and for the VPNs to work properly. DES-based encryptions are no longer supported.

How Secure Should a VPN Connection Be?

Because a VPN tunnel typically traverses a public network, most likely the Internet, you need to encrypt the connection to protect the traffic. You define the encryption and other security techniques to apply using IKE polices and IPsec proposals.

If your device license allows you to apply strong encryption, there is a wide range of encryption and hash algorithms, and Diffie-Hellman groups, from which to choose. However, as a general rule, the stronger the encryption that you apply to the tunnel, the worse the system performance. Find a balance between security and performance that provides sufficient protection without compromising efficiency.

We cannot provide specific guidance on which options to choose. If you operate within a larger corporation or other organization, there might already be defined standards that you need to meet. If not, take the time to research the options.

The following topics explain the available options.

Complying with Security Certification Requirements

Many VPN settings have options that allow you to comply with various security certification standards. Review your certification requirements and the available options to plan your VPN configuration.

Deciding Which Encryption Algorithm to Use

When deciding which encryption algorithms to use for the IKE policy or IPsec proposal, your choice is limited to algorithms supported by the devices in the VPN.

For IKEv2, you can configure multiple encryption algorithms. The system orders the settings from the most secure to the least secure and negotiates with the peer using that order. For IKEv1, you can select a single option only.

For IPsec proposals, the algorithm is used by the Encapsulating Security Protocol (ESP), which provides authentication, encryption, and anti-replay services. ESP is IP protocol type 50. In IKEv1 IPsec proposals, the algorithm name is prefixed with ESP-.

If your device license qualifies for strong encryption, you can choose from the following encryption algorithms. If you are not qualified for strong encryption, you can select DES only.

Note If you are qualified for strong encryption, before upgrading from the evaluation license to a smart license, check and update your encryption algorithms for stronger encryption so that the VPN configuration works properly. Choose AES-based algorithms. DES is not supported if you are registered using an account that supports strong encryption. After registration, you cannot deploy changes until you remove all uses of DES.

- AES-GCM—(IKEv2 only.) Advanced Encryption Standard in Galois/Counter Mode is a block cipher mode of operation providing confidentiality and data-origin authentication, and provides greater security than AES. AES-GCM offers three different key strengths: 128-, 192-, and 256-bit keys. A longer key

provides higher security but a reduction in performance. GCM is a mode of AES that is required to support NSA Suite B. NSA Suite B is a set of cryptographic algorithms that devices must support to meet federal standards for cryptographic strength. .

- AES—Advanced Encryption Standard is a symmetric cipher algorithm that provides greater security than DES and is computationally more efficient than 3DES. AES offers three different key strengths: 128-, 192-, and 256-bit keys. A longer key provides higher security but a reduction in performance.
- DES—Data Encryption Standard, which encrypts using 56-bit keys, is a symmetric secret-key block algorithm. If your license account does not meet the requirements for export controls, this is your only option.
- Null, ESP-Null—Do not use. A null encryption algorithm provides authentication without encryption. This is typically used for testing purposes only. However, it does not work at all on many platforms, including virtual.

Deciding Which Hash Algorithms to Use

In IKE policies, the hash algorithm creates a message digest, which is used to ensure message integrity. In IKEv2, the hash algorithm is separated into two options, one for the integrity algorithm, and one for the pseudo-random function (PRF).

In IPsec proposals, the hash algorithm is used by the Encapsulating Security Protocol (ESP) for authentication. In IKEv2 IPsec Proposals, this is called the integrity hash. In IKEv1 IPsec proposals, the algorithm name is prefixed with ESP-, and there is also an -HMAC suffix (which stands for "hash method authentication code").

For IKEv2, you can configure multiple hash algorithms. The system orders the settings from the most secure to the least secure and negotiates with the peer using that order. For IKEv1, you can select a single option only.

You can choose from the following hash algorithms.

- SHA (Secure Hash Algorithm)—Standard SHA (SHA1) produces a 160-bit digest.

  The following SHA-2 options, which are even more secure, are available for IKEv2 configurations. Choose one of these if you want to implement the NSA Suite B cryptography specification.

  - SHA256—Specifies the Secure Hash Algorithm SHA 2 with the 256-bit digest.
  - SHA384—Specifies the Secure Hash Algorithm SHA 2 with the 384-bit digest.
  - SHA512—Specifies the Secure Hash Algorithm SHA 2 with the 512-bit digest.

- Null or None (NULL, ESP-NONE)—(IPsec Proposals only.) A null Hash Algorithm; this is typically used for testing purposes only. However, you should choose the null integrity algorithm if you select one of the AES-GCM options as the encryption algorithm. Even if you choose a non-null option, the integrity hash is ignored for these encryption standards.

Deciding Which Diffie-Hellman Modulus Group to Use

You can use the following Diffie-Hellman key derivation algorithms to generate IPsec security association (SA) keys. Each group has a different size modulus. A larger modulus provides higher security, but requires more processing time. You must have a matching modulus group on both peers.

If you select AES encryption, to support the large key sizes required by AES, you should use Diffie-Hellman (DH) Group 5 or higher. IKEv1 policies do not support all of the groups listed below.

To implement the NSA Suite B cryptography specification, use IKEv2 and select one of the elliptic curve Diffie-Hellman (ECDH) options: 19, 20, or 21. Elliptic curve options and groups that use 2048-bit modulus are less exposed to attacks such as Logjam.

For IKEv2, you can configure multiple groups. The system orders the settings from the most secure to the least secure and negotiates with the peer using that order. For IKEv1, you can select a single option only.

- 14—Diffie-Hellman Group 14: 2048-bit modular exponential (MODP) group. Considered good protection for 192-bit keys.
- 15—Diffie-Hellman Group 15: 3072-bit MODP group.
- 16—Diffie-Hellman Group 16: 4096-bit MODP group.
- 19—Diffie-Hellman Group 19: National Institute of Standards and Technology (NIST) 256-bit elliptic curve modulo a prime (ECP) group.
- 20—Diffie-Hellman Group 20: NIST 384-bit ECP group.
- 21—Diffie-Hellman Group 21: NIST 521-bit ECP group.
- 31—Diffie-Hellman Group 31: Curve25519 256-bit EC Group.

Deciding Which Authentication Method to Use

Preshared keys and digital certificates are the methods of authentication available for VPNs.

Site-to-site, IKEv1 and IKEv2 VPN connections can use both options.

Remote Access, which uses SSL and IPsec IKEv2 only, supports digital certificate authentication only.

Preshared keys allow for a secret key to be shared between two peers and used by IKE during the authentication phase. The same shared key must be configured at each peer or the IKE SA cannot be established.

Digital certificates use RSA key pairs to sign and encrypt IKE key management messages. Certificates provide non-repudiation of communication between two peers, meaning that it can be proved that the communication actually took place. When using this authentication method, you need a Public Key Infrastructure (PKI) defined where peers can obtain digital certificates from a Certification Authority (CA). CAs manage certificate requests and issue certificates to participating network devices providing centralized key management for all of the participating devices.

Preshared keys do not scale well, using a CA improves the manageability and scalability of your IPsec network. With a CA, you do not need to configure keys between all encrypting devices. Instead, each participating device is registered with the CA, and requests a certificate from the CA. Each device that has its own certificate and the public key of the CA can authenticate every other device within a given CA's domain.

Pre-shared Keys

Pre-shared key enables you to share a secret key between two peers. IKE uses the key in the authentication phase. You must configure the same shared key on each peer, or the IKE SA cannot be established.

To configure the pre-shared keys, choose whether you want to use a manual or automatically generated key, and then specify the key in the IKEv1/IKEv2 options. Then, when you deploy your configuration, the key is configured on all devices in the topology.

PKI Infrastructure and Digital Certificates

Public Key Infrastructure

A PKI provides centralized key management for participating network devices. It is a defined set of policies, procedures, and roles that support *public key cryptography* by generating, verifying, and revoking *public key certificates* commonly known as *digital certificates*.

In public key cryptography, each endpoint of a connection has a key pair consisting of both a public and a private key. The key pairs are used by the VPN endpoints to sign and encrypt messages. The keys act as complements, and anything encrypted with one of the keys can be decrypted with the other, securing the data flowing over the connection.

Generate a general purpose RSA, ECDSA, or EDDSA key pair, used for both signing and encryption, or you generate separate key pairs for each purpose. Separate signing and encryption keys help to reduce exposure of the keys. SSL uses a key for encryption but not signing, however, IKE uses a key for signing but not encryption. By using separate keys for each, exposure of the keys is minimized.

Digital Certificates or Identify Certificates

When you use Digital Certificates as the authentication method for VPN connections, peers are configured to obtain digital certificates from a Certificate Authority (CA). CAs are trusted authorities that “sign” certificates to verify their authenticity, thereby guaranteeing the identity of the device or user.

CA servers manage public CA certificate requests and issue certificates to participating network devices as part of a Public Key Infrastructure (PKI), this activity is called Certificate Enrollment. These digital certificates, also called identity certificates contain:

- The digital identification of the owner for authentication, such as name, serial number, company, department, or IP address.
- A public key needed to send and receive encrypted data to the certificate owner.
- The secure digital signature of a CA.

Certificates also provide non-repudiation of communication between two peers, meaning that it they prove that the communication actually took place.

Certificate Enrollment

Using a PKI improves the manageability and scalability of your VPN since you do not have to configure pre-shared keys between all the encrypting devices. Instead, you individually *enroll* each participating device with a CA server, which is explicitly trusted to validate identities and create an identity certificate for the device. When this has been accomplished, each participating peer sends their identity certificate to the other peer to validate their identities and establish encrypted sessions with the public keys contained in the certificates. See Certificate Enrollment Objects, on page 1413 for details on enrolling Firewall Threat Defense devices.

Certificate Authority Certificates

In order to validate a peer’s certificate, each participating device must retrieve the CA's certificate from the server. A CA certificate is used to sign other certificates. It is self-signed and called a root certificate. This certificate contains the public key of the CA, used to decrypt and validate the CA's digital signature and the contents of the received peer's certificate. The CA certificate may be obtained by:

- Using the Simple Certificate Enrollment Protocol (SCEP) or Enrollment over Secure Transport (EST) to retrieve the CA's certificate from the CA server
- Manually copying the CA's certificate from another participating device

Trustpoints

Once enrollment is complete, a trustpoint is created on the managed device. It is the object representation of a CA and associated certificates. A trustpoint includes the identity of the CA, CA-specific parameters, and an association with a single enrolled identity certificate.

PKCS#12 File

A PKCS#12, or PFX, file holds the server certificate, any intermediate certificates, and the private key in one encrypted file. This type of file may be imported directly into a device to create a trustpoint.

Revocation Checking

A CA may also revoke certificates for peers that no longer participate in you network. Revoked certificates are either managed by an Online Certificate Status Protocol (OCSP) server or are listed in a certificate revocation list (CRL) stored on an LDAP server. A peer may check these before accepting a certificate from another peer.

Removed or Deprecated Hash Algorithms, Encryption Algorithms, and Diffie-Hellman Modulus Groups

Support has been removed for less secure ciphers. We recommend that you update your VPN configuration before you upgrade to Firewall Threat Defense 6.70 to supported DH and encryption algorithms to ensure the VPN works correctly.

Update your IKE proposals and IPSec policies to match the ones supported in Firewall Threat Defense 6.70 and then deploy the configuration changes.

The following less secure ciphers have been removed or deprecated in Firewall Threat Defense 6.70 onwards:

- **Diffie-Hellman GROUP 5** is deprecated for IKEv1 and IKEv2.
- Diffie-Hellman groups 2 and 24 have been removed.
- **Encryption algorithms**: 3DES, AES-GMAC, AES-GMAC-192, AES-GMAC-256 have been removed.

Note **DES** continues to be supported in evaluation mode or for users who do not satisfy export controls for strong encryption.

NULL is removed in IKEv2 policy, but supported in both IKEv1 and IKEv2 IPsec transform-sets.

VPN Topology Options

When you create a new VPN topology you must, at minimum, give it a unique name, specify a topology type, and select the IKE version. You can select from three types of topologies, each containing a group of VPN tunnels:

- Point-to-point (PTP) topologies establish a VPN tunnel between two endpoints.
- Hub and Spoke topologies establish a group of VPN tunnels connecting a hub endpoint to a group of spoke endpoints.
- Full Mesh topologies establish a group of VPN tunnels among a set of endpoints.

Define a pre-shared key for VPN authentication manually or automatically, there is no default key. When choosing automatic, the Secure Firewall Management Center generates a pre-shared key and assigns it to all the nodes in the topology.

Point-to-Point VPN Topology

In a point-to-point VPN topology, two endpoints communicate directly with each other. You configure the two endpoints as peer devices, and either device can start the secured connection.

The following diagram displays a typical point-to-point VPN topology.

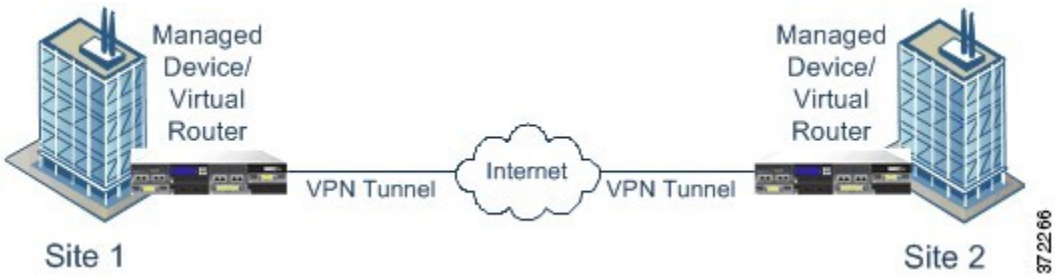

Hub and Spoke VPN Topology

In a Hub and Spoke VPN topology, a central endpoint (hub node) connects with multiple remote endpoints (spoke nodes). Each connection between the hub node and an individual spoke endpoint is a separate VPN tunnel. The hosts behind any of the spoke nodes can communicate with each other through the hub node.

The Hub and Spoke topology commonly represent a VPN that connects an organization's main and branch office locations using secure connections over the Internet or other third-party network. These deployments provide all employees with controlled access to the organization's network. Typically, the hub node is located at the main office. Spoke nodes are located at branch offices and start most of the traffic.

The following diagram displays a typical Hub and Spoke VPN topology.

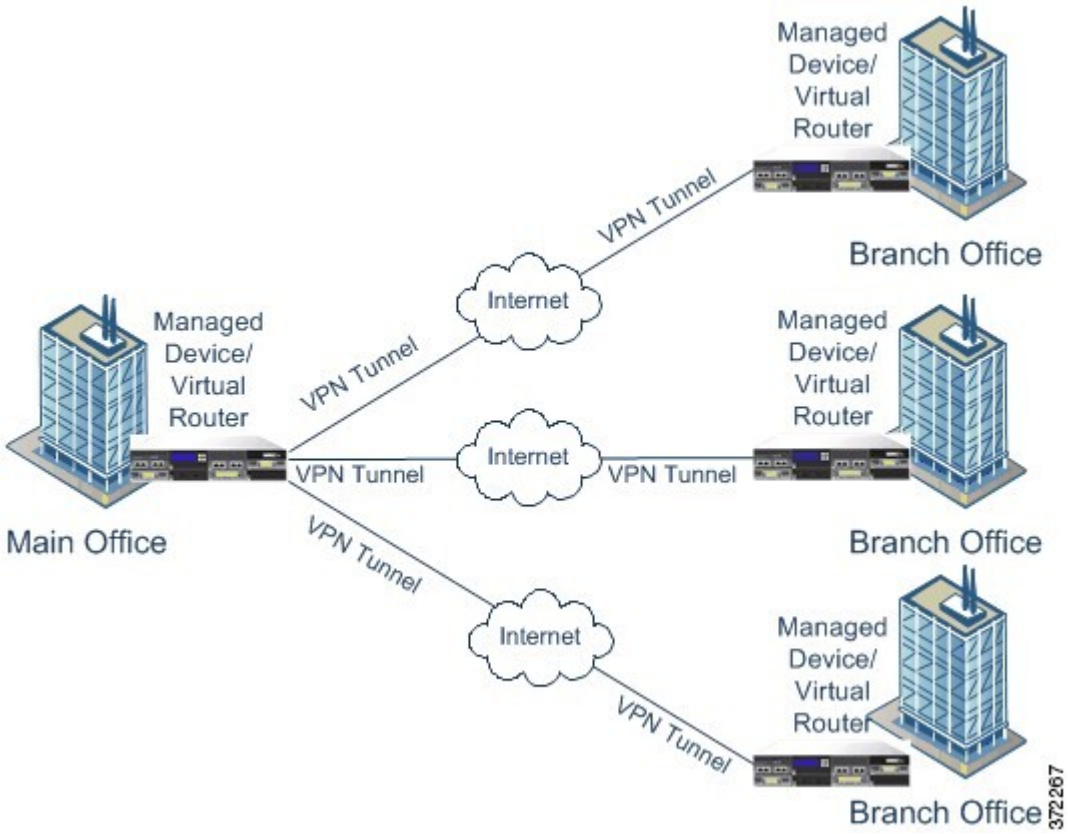

Full Mesh VPN Topology

In a Full Mesh VPN topology, all endpoints can communicate with every other endpoint by an individual VPN tunnel. This topology offers redundancy so that when one endpoint fails, the remaining endpoints can still communicate with each other. It commonly represents a VPN that connects a group of decentralized branch office locations. The number of VPN-enabled managed devices you deploy in this configuration depends on the level of redundancy you require.

The following diagram displays a typical Full Mesh VPN topology.

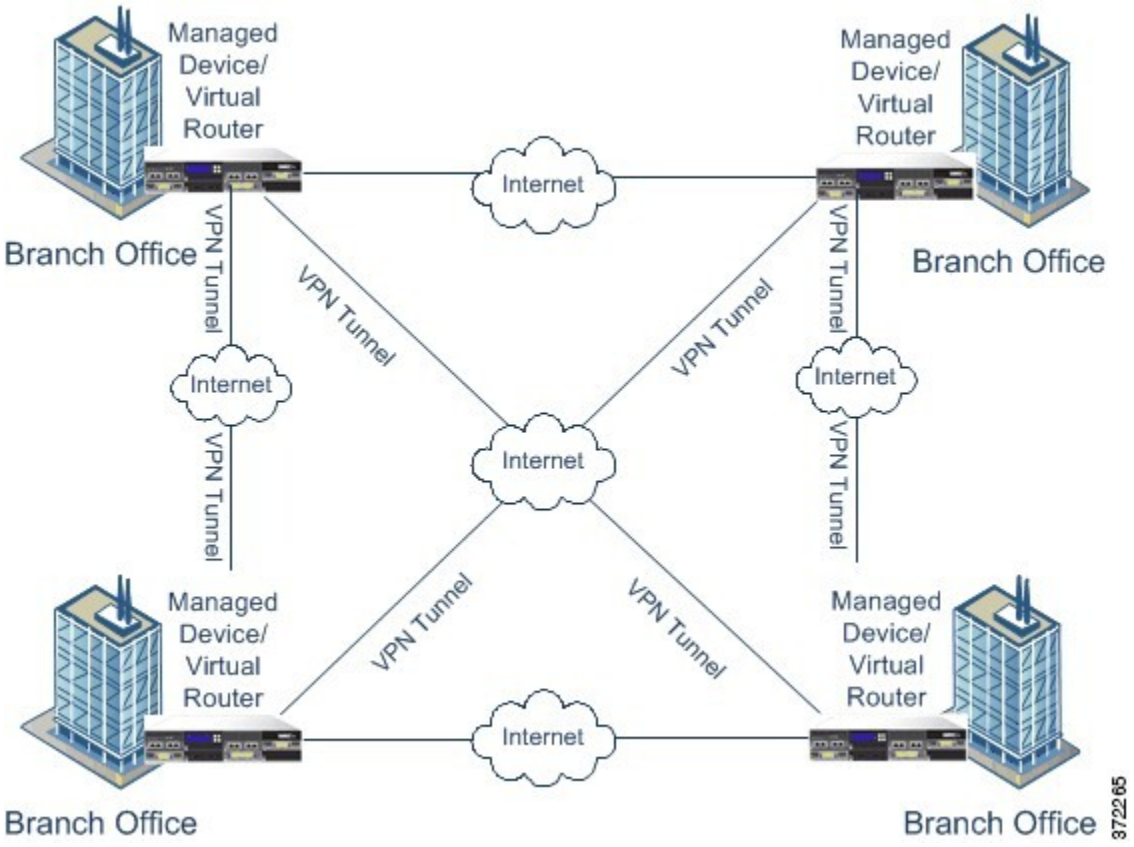

Implicit Topologies

In addition to the three main VPN topologies, other more complex topologies can be created as combinations of these topologies. They include:

- Partial mesh—A network in which some devices are organized in a full mesh topology, and other devices form either a hub-and-spoke or a point-to-point connection to some of the fully meshed devices. A partial mesh does not provide the level of redundancy of a full mesh topology, but it is less expensive to implement. Partial mesh topologies are used in peripheral networks that connect to a fully meshed backbone.
- Tiered hub-and-spoke—A network of hub-and-spoke topologies in which a device can behave as a hub in one or more topologies and a spoke in other topologies. Traffic is permitted from spoke groups to their most immediate hub.
- Joined hub-and-spoke—A combination of two topologies (hub-and-spoke, point-to-point, or full mesh) that connect to form a point-to-point tunnel. For example, a joined hub-and-spoke topology could comprise two hub-and-spoke topologies, with the hubs acting as peer devices in a point-to-point topology.

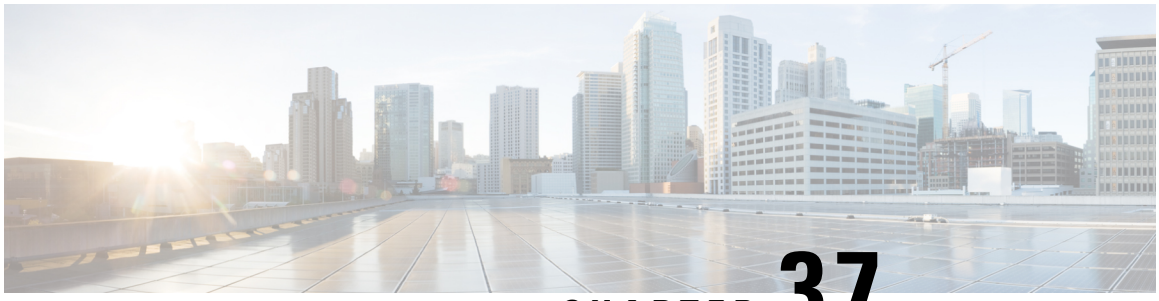

CHAPTER 37

Site-to-Site VPNs

- About Site-to-Site VPN, on page 1511
- Types of Site-to-Site VPN Topologies, on page 1514
- Requirements and Prerequisites for Site-to-Site VPN, on page 1514
- Manage Site-to-Site VPNs, on page 1515
- Configure a Policy-based Site-to-Site VPN, on page 1516
- About Virtual Tunnel Interfaces, on page 1529
- Guidelines and Limitations for Virtual Tunnel Interfaces, on page 1533
- Add a VTI Interface, on page 1537
- Create a Route-based Site-to-Site VPN, on page 1538
- Route Traffic Through a Backup VTI Tunnel, on page 1549
- Configure Dynamic VTI for a Route-based Site-to-Site VPN, on page 1551
- How to Configure a Virtual Router with Dynamic VTI, on page 1552
- Configure Routing and AC Policies for VTI, on page 1552
- View Virtual Tunnel Information, on page 1555
- Deploy a SASE Tunnel on Umbrella, on page 1555
- Guidelines and Limitations for Configuring SASE Tunnels on Umbrella, on page 1556
- How to Deploy a SASE Tunnel on Umbrella, on page 1557
- Monitoring the Site-to-Site VPNs, on page 1562
- History for Site-to-Site VPN, on page 1569

About Site-to-Site VPN

Secure Firewall Threat Defense site-to-site VPN supports the following features:

- Both IPsec IKEv1 & IKEv2 protocols.
- Certificates and automatic or manual preshared keys for authentication.
- IPv4 & IPv6. All combinations of inside and outside are supported.
- IPsec IKEv2 Site-to-Site VPN topologies provide configuration settings to comply with security certifications.
- Static and dynamic Interfaces.
- HA environments for both Firewall Management Center and Firewall Threat Defense.

- VPN alerts when the tunnel goes down.
- Tunnel statistics available using the Firewall Threat Defense Unified CLI.
- IKEv1 and IKEv2 back-up peer configuration for point-to-point extranet and hub-and-spoke VPNs.
- Extranet device as hub in 'Hub and Spokes' deployments.
- Dynamic IP address for a managed endpoint pairing with extranet device in 'Point to Point' deployments.
- Dynamic IP address for extranet device as an endpoint.
- Hub as extranet in 'Hub and Spokes' deployments.

VPN Topology

To create a new site-to-site VPN topology you must, specify a unique name, a topology type, choose the IKE version that is used for IPsec IKEv1 or IKEv2, or both. Also, determine your authentication method. Once configured, you deploy the topology to Firewall Threat Defense devices. The Secure Firewall Management Center configures site-to-site VPNs on Firewall Threat Defense devices only.

You can select from three types of topologies, containing one or more VPN tunnels:

- Point-to-point (PTP) deployments establish a VPN tunnel between two endpoints.
- Hub and Spoke deployments establish a group of VPN tunnels connecting a hub endpoint to a group of spoke nodes.
- Full Mesh deployments establish a group of VPN tunnels among a set of endpoints.

IPsec and IKE

In the Secure Firewall Management Center, site-to-site VPNs are configured based on IKE policies and IPsec proposals that are assigned to VPN topologies. Policies and proposals are sets of parameters that define the characteristics of a site-to-site VPN, such as the security protocols and algorithms that are used to secure traffic in an IPsec tunnel. Several policy types may be required to define a full configuration image that can be assigned to a VPN topology.

Authentication

For authentication of VPN connections, configure a preshared key in the topology, or a trustpoint on each device. Preshared keys allow for a secret key, used during the IKE authentication phase, to be shared between two peers. A trustpoint includes the identity of the CA, CA-specific parameters, and an association with a single enrolled identity certificate.

Extranet Devices

Each topology type can include extranet devices, devices that you don’t manage in Firewall Management Center. These include:

- Cisco devices that Secure Firewall Management Center supports, but for which your organization isn’t responsible. Such as spokes in networks managed by other organizations within your company, or a connection to a service provider or partner's network.
- Non-Cisco devices. You can’t use Secure Firewall Management Center to create and deploy configurations to non-Cisco devices.

Add non-Cisco devices, or Cisco devices not managed by the Secure Firewall Management Center, to a VPN topology as "Extranet" devices. Also specify the IP address of each remote device.

Secure Firewall Threat Defense Site-to-site VPN Guidelines and Limitations

- Site-to-site VPN supports ECMP zone interfaces.
- You must configure all nodes in a topology with either crypto ACL or a protected network. You cannot configure a topology with crypto ACL on one node and protected network on another.
- You can configure a VPN connection across domains by using an extranet peer for the endpoint not in the current domain.
- You can backup Firewall Threat Defense VPNs using the Firewall Management Center backup.
- IKEv1 does not support CC/UCAPL-compliant devices. We recommend that you use IKEv2 for these devices.
- You cannot move a VPN topology between domains.
- VPN does not support network objects with a 'range' option.
- Firewall Threat Defense VPNs do not currently support PDF export and policy comparison.
- There is no per-tunnel or per-device edit option for Firewall Threat Defense VPNs, you can edit only the whole topology.
- The Firewall Management Center does not verify the device interface address verification for transport mode when you select a crypto ACL.
- There is no support for automatic mirror ACE generation. Mirror ACE generation for the peer is a manual process on either side.
- With crypto ACL, the Firewall Management Center supports only point to point VPN and does not support tunnel health events.
- Whenever IKE ports 500/4500 are in use or when there are some active PAT translations, you cannot configure a site-to-site VPN on the same ports as it fails to start the service on those ports.
- Tunnel status is not updated in realtime, but at an interval of five minutes in the Firewall Management Center.
- You cannot use the character " (double quote) as part of pre-shared keys. If you have used " in a pre-shared key, ensure that you change the character.
- In a site-to-site VPN configuration with two devices managed by the same Firewall Management Center, you cannot configure the devices as backup peers. You must configure one of peer devices in the topology as an extranet device.
- Configure unique local IKE identity for all tunnels across all your VPN topologies.
- In a remote branch deployment (RBD) with High Availability when you break a High Availability pair:
    - For Hub and Spoke VPN:
        - If the Firewall Threat Defense HA is a hub, VPN configurations will be available in the active device and will be removed in the standby device.

- If the Firewall Threat Defense HA is a spoke, VPN configurations will be available in the active and standby devices.

- For Point-to-point VPN, if the endpoint has an RBD Firewall Threat Defense HA WAN interface, VPN configurations will be removed from the standby device.

Types of Site-to-Site VPN Topologies

Site-to-Site VPN Topology	Description	More Information
SD-WAN Topology	Configure a secure branch network using the SD-WAN wizard. This wizard simplifies and automates VPN and routing configuration of your network using a hub and spoke topology, enabling SD-WAN capabilities.	Using SD-WAN Wizard for Secure Branch Network Deployment, on page 2013
Route-Based VPN	Configure secure traffic dynamically between peers in a network based on routing over Virtual Tunnel Interfaces (VTIs).	Create a Route-based Site-to-Site VPN, on page 1538
Policy-Based VPN	Configure secure traffic between peers in a network based on a static policy using protected networks.	Configure a Policy-based Site-to-Site VPN, on page 1516
SASE Topology	Configure an IPsec IKEv2 tunnel from a threat defense device to an Umbrella Secure Internet Gateway (SIG). This tunnel forwards all internet-bound traffic to the Umbrella SIG for inspection and filtering.	Configure a SASE Tunnel for Umbrella, on page 1560

Requirements and Prerequisites for Site-to-Site VPN

Model Support

Threat Defense

Supported Domains

Leaf

User Roles

Admin

Supported Interfaces

Topology Type	Interface Type

Policy-Based and SD-WAN	• Physical interfaces
	  • Non-management
	  • Interface Mode must be either Routed or None
	• Subinterface interfaces
	• Redundant interfaces
	• Etherchannel interfaces
	• VLAN interfaces

Route-Based	Static Virtual Tunnel Interfaces

Manage Site-to-Site VPNs

The Site-to-Site VPN page provides a snapshot of site-to-site VPN tunnels. You can view the status of the tunnels and filter the tunnels based on the device, topology, or tunnel type. The page lists 20 topologies per page and you can navigate between pages to view more topology details. You can click individual VPN topologies to expand and view details of the endpoints.

Before you begin

For certificate authentication of your site-to-site VPNs, you must prepare the devices by allocating trustpoints as described in Certificates, on page 1485.

Procedure

Select **Devices > VPN > Site To Site** to manage your Firewall Threat Defense site-to-site VPN configurations and deployments.

The page lists the site-to-site VPNs topologies and indicates the status of tunnels using color codes:

- Active (Green)–There is an active IPsec Tunnel.
- Unknown (Amber)–No tunnel establishment event has been received from the device yet.
- Down (Red)–There are no active IPsec tunnels.
- Deployment Pending–Topology has not been deployed on the device yet.

Choose from the following:

- **Refresh**—View the updated status of the VPNs.
- **Add**—Create new policy based or route-based Site to Site VPNs.

- **Edit**—Modify the settings of an existing VPN topology.

  Note
  You cannot edit the topology type after you initially save it. To change the topology type, delete the topology and create a new one.

  Two users shouldn't edit the same topology simultaneously; however, the web interface doesn't prevent simultaneous editing.

- **Delete**—To delete a VPN deployment, click **Delete** ( ).
- Deploy—Choose **Deploy** > **Deployment**; see Deploy Configuration Changes, on page 259.

  Note
  Some VPN settings are validated only during deployment. Be sure to verify that your deployment was successful.

Configure a Policy-based Site-to-Site VPN

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site**, and click **Add**.

Step 2 Enter a name for the topology in the **Topology Name** field.

Step 3 Click the **Policy-Based VPN** radio button.

Step 4 Select the VPN topology and click **Create**.

Step 5 Choose the IKE versions to use during IKE negotiations. Check the **IKEv1** or **IKEv2** check box.

Default is IKEv2. Select either or both options as appropriate; select IKEv1 if any device in the topology doesn't support IKEv2.

You can also configure a backup peer for point-to-point extranet VPNs. For more information, see Firewall Threat Defense VPN Endpoint Options, on page 1517.

Step 6 Required: Add Endpoints for this VPN deployment by clicking **Add** ( ) for each node in the topology.

Configure each endpoint field as described in Firewall Threat Defense VPN Endpoint Options, on page 1517.

- For Point to point, configure **Node A** and **Node B**.
- For Hub and Spoke, configure a **Hub Node** and **Spoke Nodes**
- For Full Mesh, configure multiple **Nodes**

Step 7 (Optional) Specify non-default IKE options for this deployment as described in Firewall Threat Defense VPN IKE Options, on page 1521

Step 8 (Optional) Specify non-default IPsec options for this deployment as described in Firewall Threat Defense VPN IPsec Options, on page 1523

Step 9 (Optional) Specify non-default Advanced options for this deployment as described in Firewall Threat Defense Advanced Site-to-site VPN Deployment Options, on page 1526.

Step 10 Click **Save**.
The endpoints are added to your configuration.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Note Some VPN settings are validated only during deployment. Be sure to verify that your deployment was successful.

If you get an alert that your VPN tunnel is inactive even when the VPN session is up, follow the VPN troubleshooting instructions to verify and ensure that your VPN is active. For more information, see VPN Monitoring and Troubleshooting, on page 1703 and VPN Troubleshooting, on page 1712.

Firewall Threat Defense VPN Endpoint Options

Navigation Path

Configure the basic parameters for a point-to-point topology in a route-based VPN as described in Configure Endpoints for a Point to Point Topology and click the **Endpoints** tab.

Fields

Device

Choose an endpoint node for your deployment:

- A Firewall Threat Defense device managed by this Firewall Management Center
- A Firewall Threat Defense high availability container managed by this Firewall Management Center
- An **Extranet** device, any device (Cisco or third party) not managed by this Firewall Management Center.

Device Name

For extranet devices only, provide a name for this device. We recommend naming it such that it is identifiable as an unmanaged device.

Interface

If you chose a managed device as your endpoint, choose an interface on that managed device.

For 'Point to Point' deployments, you can also configure an endpoint with dynamic interface. An endpoint with a dynamic interface can pair only with an extranet device and can’t pair with an endpoint, which has a managed device.

You can configure device interfaces at **Devices** > **Device ManagementAdd/Edit device** > **Interfaces**.

IP Address

- If you choose an extranet device, a device **not** managed by the Firewall Management Center, specify an IP address for the endpoint.

  For an extranet device, select **Static** and specify an IP address or select **Dynamic** to allow dynamic extranet devices.

- If you chose a managed device as an endpoint, choose a single IPv4 address or multiple IPv6 addresses from the drop-down list. These IP addresses are already assigned to this interface on the managed device.
- All endpoints in a topology must have the same IP addressing scheme. IPv4 tunnels can carry IPv6 traffic and vice versa. The Protected Networks define which addressing scheme the tunneled traffic uses.
- If the managed device is a high-availability container, choose from a list of interfaces.

This IP is Private

Check the check box if the endpoint resides behind a firewall with network address translation (NAT).

Note Use this option only when the peer is managed by the same Firewall Management Center and don't use this option if the peer is an extranet device.

Public IP address

If you checked the **This IP is Private** check box, specify a public IP address for the firewall. If the endpoint is a responder, specify this value.

Connection Type

Specify the allowed negotiation as bidirectional, answer-only, or originate-only. Supported combinations for the connection type are:

Table 78: Connection Type Supported Combinations

Remote Node	Central Node
Originate-Only	Answer-Only
Bi-Directional	Answer-Only
Bi-Directional	Bi-Directional

Certificate Map

Choose a preconfigured certificate map object, or click **Add** (+) to add a certificate map object. The certificate map defines what information is necessary in the received client certificate to be valid for VPN connectivity. See Certificate Map Objects, on page 1462 for details.

Protected Networks

Caution Hub and Spoke topology—To avoid traffic drop for a dynamic crypto map, ensure that you don't select the protected network *any* for both the endpoints.

If the protected network is configured as *any*, on both the endpoints, the crypto ACL that works upon the tunnel is not generated.

Defines the networks that are protected by this VPN endpoint. Select the networks by selecting the list of Subnet/IP Address that define the networks that are protected by this endpoint. Click **Add** (+) to select from available Network Objects or add new Network Objects. See Creating Network Objects, on page 1403. Access control lists are generated from the choices made here.

- **Subnet/IP Address (Network)**—VPN endpoints can't have the same IP address and protected networks in a VPN endpoint pair cannot overlap. If protected networks for an endpoint contain IPv4 or IPv6 entries, the other endpoint's protected network must have at least one entry of the same type (IPv4 or IPv6). If it doesn't, the other endpoint's IP address must be of the same type and not overlap with the entries in the protected network. (Use /32 CIDR address blocks for IPv4 and /128 CIDR address blocks for IPv6.) If both of these checks fail, the endpoint pair is invalid.

Note By default, **Reverse Route Injection is enabled** is enabled in Firewall Management Center.

Subnet/IP Address (Network) remains the default selection.

When you've selected Protected Networks as *Any* and observe default route traffic being dropped, disable the Reverse Route Injection. Choose **VPNSite to Site** > edit a VPN > **IPsec** > **Enable Reverse Route Injection**. Deploy the configuration changes to remove set reverse-route (Reverse Route Injection) from the crypto map configuration and remove the VPN-advertised reverse route that causes the reverse tunnel traffic to be dropped.

- **Access List (Extended)**—An extended access list provides the capability to control the type of traffic that will be accepted by this endpoint, like GRE or OSPF traffic. Traffic may be restricted either by address or port. Click **Add** (+) to add access control list objects.

Note Access Control List is supported only in the point to point topology.

Enable NAT Traversal

Check this check box to allow seamless communication between the peer Firewall Threat Defense devices when there are NAT devices between these devices. For hub and spoke topologies, this option is available only for the spoke. You can uncheck this check box to disable this feature for an endpoint. This parameter is a per peer setting for an endpoint within a topology.

To view or configure the global NAT Traversal (NAT-T) setting, in the **Add/Edit VPN Topology** dialog box:

1. Click the **Advanced** tab.
2. In the left navigation pane, click **Tunnel**.
3. Under **NAT Settings**, the **Keepalive Messages Traversal** is a global setting that enables NAT-T for all endpoints within a topology.

Exempt VPN traffic from network address translation

Check this check box to exempt the VPN traffic from the Network Address Translation (NAT) rules.

If you do not exempt the VPN traffic from the NAT rules, the traffic gets dropped or is not routed through the VPN tunnel to the remote device. After you enable this option, you can view the NAT exemptions for the device in the NAT policy page (**Device > NAT > NAT Exemptions**).

Inside interfaces directly connected to the internal network

Specify the security zone or interface group for the inside interface(s) where the protected networks reside. By default, the inside interface is any.

Click **+** to configure one or more interfaces from a security zone or an interface group that can map to one or more inside interfaces. Ensure that the interface type of the security zone or an interface group is Routed.

Advanced Settings

Enable Dynamic Reverse Route Injection—Reverse Route Injection (RRI) enables routes to be automatically inserted into the routing process, for the networks and hosts protected by a remote tunnel endpoint. Dynamic RRI routes are created only upon the successful establishment of IPsec security associations (SA's).

Note

- Dynamic RRI is supported only on IKEv2, and not supported on IKEv1 or IKEv1 + IKEv2.
- Dynamic RRI isn't supported on originate-only peer, Full Mesh topology, and Extranet peer.
- In Point-to-Point, only one peer can have dynamic RRI enabled.
- Between Hub and Spoke, only one of the endpoints can have dynamic RRI enabled.
- Dynamic RRI cannot be combined with a dynamic crypto map.

Send Local Identity to Peers—Select this option to send local identity information to the peer device. Select one of the following **Local Identity Configuration** from the list and configure the local identity:

- **IP address**—Use the IP address of the interface for the identity.
- **Auto**—Use the IP address for pre-shared key and Cert DN for certificate-based connections.
- **Email ID**—Specify the email ID to use for the identity. The email ID can be up to 127 characters.
- **Hostname**—Use the fully qualified hostname.
- **Key ID**—Specify the key-id to use for the identity. The key ID must be fewer than 65 characters.

The local identity is used to configure a unique identity per IKEv2 tunnel, instead of a global identity for all the tunnels. The unique identity allows Firewall Threat Defense to have multiple IPsec tunnels behind a NAT to connect to the Cisco Umbrella Secure Internet Gateway (SIG).

For information about configuring a unique tunnel ID on Umbrella, see **Cisco Umbrella SIG User Guide**.

VPN Filter—Select an extended access list from the list or click **Add** to create a new extended access list object to filter the site-to-site VPN traffic.

The VPN filter provides more security and filters site-to-site VPN data using an extended access list. The extended access list object selected for the VPN filter lets you filter pre-encrypted traffic before entering the VPN tunnel and decrypted traffic that exits a VPN tunnel. The **sysopt permit-vpn** option, when enabled, would bypass the access control policy rules for the traffic coming from the VPN tunnel. When the **sysopt permit-vpn** option is enabled, the VPN filter helps in identifying and filtering the site-to-site VPN traffic.

Note The VPN filter is supported only on Point to Point, and Hub and Spoke topologies. It isn't supported on Mesh topology.

For Hub and Spoke topology, you can choose to override the hub VPN filter on the spoke endpoints in case a different VPN filter needs to enabled on a specific tunnel.

Select the **Override VPN Filter on the Hub** option to override the hub VPN filter on the spokes. Select the **Remote VPN Filter** extended access list object or create an access list to override.

Note For an extranet device as a spoke, only the **Override VPN filter on the Hub** option is available.

For more information about sysopt permit-VPN, see Firewall Threat Defense Advanced Site-to-site VPN Tunnel Options, on page 1528.

Firewall Threat Defense VPN IKE Options

For the versions of IKE you have chosen for this topology, specify the **IKEv1/IKEv2 Settings**.

Note Settings in this dialog apply to the entire topology, all tunnels, and all managed devices.

Navigation Path

Configure the basic parameters for a point-to-point topology in a route-based VPN as described in Configure Endpoints for a Point to Point Topology and click the **IKE** tab.

Fields

Policy

Choose the required IKEv1 or IKEv2 policy objects from the predefined list or create new objects to use. You can choose multiple IKEv1 and IKEv2 policies. IKEv1 and IKEv2 support a maximum of 20 IKE policies, each with a different set of values. Assign a unique priority to each policy that you create. The lower the priority number, the higher the priority.

We recommend that you do not use values such as 10, 20 and so on because the default IKEv2 policy for remote access VPN can have these values as its priority value. Before deployment, verify that the priority values of the IKE policies (site-to-site and remote access VPN) do not conflict.
For details, see Firewall Threat Defense IKE Policies, on page 1475

Authentication Type

Site-to-site VPN supports two authentication methods, pre-shared key and certificate. For an explanation of the two methods, see Deciding Which Authentication Method to Use, on page 1504.

Note In a VPN topology that supports IKEv1, the **Authentication Method** specified in the chosen IKEv1 Policy object becomes the default in the IKEv1 **Authentication Type** setting. These values must match, otherwise, your configuration will error.

- **Pre-shared Automatic Key**—The Firewall Management Center automatically defines the pre-shared key for this VPN. Specify the **Pre-shared Key Length**, the number of characters in the key, 1-27.

  The character " (double quote) isn't supported as part of pre-shared keys. If you've used " in a pre-shared key, ensure that you change the character after you upgrade to Secure Firewall Threat Defense 6.30 or higher.

- **Pre-shared Manual Key**—Manually assign the pre-shared key for this VPN. Specify the **Key** and then reenter the same to **Confirm Key**.

  When you choose this option for IKEv2, the **Enforce hex-based pre-shared key only** check box appears, check if desired. If enforced, you must enter a valid hex value for the key, an even number of 2-256 characters, using numerals 0-9, or A-F.

- **Certificate**—When you use certificates as the authentication method for VPN connections, peers obtain digital certificates from a CA server in your PKI infrastructure, and trade them to authenticate each other.

  In the **Certificate** field, select a preconfigured certificate enrollment object. This enrollment object generates a trustpoint with the same name on the managed device. The certificate enrollment object should be associated with and installed on the device, post which the enrollment process is complete, and then a trustpoint is created.

  A trustpoint is a representation of a CA or identity pair. A trustpoint includes the identity of the CA, CA-specific configuration parameters, and an association with one enrolled identity certificate.

  Before you select this option, note the following:

  - Ensure you've enrolled a certificate enrollment object on all the endpoints in the topology—A certificate enrollment object contains the Certification Authority (CA) server information and enrollment parameters that are required for creating Certificate Signing Requests (CSRs) and obtaining Identity Certificates from the specified CA. Certificate Enrollment Objects are used to enroll your managed devices into your PKI infrastructure, and create trustpoints (CA objects) on devices that support VPN connections. For instructions on creating a certificate enrollment object, see Adding Certificate Enrollment Objects, on page 1414, and for instructions on enrolling the object on the endpoints see one of the following as applicable:

    - Installing a Certificate Using Self-Signed Enrollment , on page 1489
    - Installing a Certificate using EST Enrollment, on page 1490
    - Installing a Certificate Using SCEP Enrollment, on page 1491

- Installing a Certificate Using Manual Enrollment, on page 1491
- Installing a Certificate Using a PKCS12 File, on page 1492

Note For a site-to-site VPN topology, ensure that the same certificate enrollment object is enrolled in all the endpoints in the topology. For further details, see the table below.

- Refer the following table to understand the enrollment requirement for different scenarios. Some of the scenarios require you to override the certificate enrollment object for specific devices. See Managing Object Overrides, on page 1362 to understand how to override objects.

Certificate Enrollment Types	Device identity certificate for all endpoints is from the same CA		Device identity certificate for all endpoints is from different CAs
	Device-specific parameters are NOT specified in the certificate enrollment object	Device-specific parameters are specified in the certificate enrollment object	
Manual	No override required	Override required	Override required
EST	No override required	Override required	Override required
SCEP	No override required	Override required	Override required
PKCS	Override required	Override required	Override required
Self-signed	Not applicable	Not applicable	Not applicable

- Understand the VPN certificate limitations mentioned in Secure Firewall Threat Defense VPN Certificate Guidelines and Limitations, on page 1485.

Note If you use a Windows Certificate Authority (CA), the default application policies extension is **IP security IKE intermediate**. If you use this default setting, you must select the **Ignore IPsec Key Usage** option in the Advanced Settings section on the **Key** tab in the **PKI Certificate Enrollment** dialog box for the object you select. Otherwise, the endpoints can't complete the site-to-site VPN connection.

Firewall Threat Defense VPN IPsec Options

Note Settings in this dialog apply to the entire topology, all tunnels, and all managed devices.

Configure the basic parameters for a point-to-point topology in a route-based VPN as described in Configure Endpoints for a Point to Point Topology and click the **IPsec** tab.

Crypto-Map Type

A crypto map combines all the components required to set up IPsec security associations (SA). When two peers try to establish an SA, they must each have at least one compatible crypto map entry. The IPsec security negotiation uses the proposals defined in the crypto map entry to protect the data flows specified by that crypto map's IPsec rules. Choose static or dynamic for this deployment's crypto-map:

- **Static**—Use a static crypto map in a point-to-point or full mesh VPN topology.
- **Dynamic**—Dynamic crypto-maps essentially create a crypto map entry without all the parameters configured. The missing parameters are later dynamically configured (as the result of an IPsec negotiation) to match a remote peer's requirements.

  Dynamic crypto map policies are applicable to both hub-and-spoke and point-to-point VPN topologies. To apply these policies, specify a dynamic IP address for one of the peers in the topology and ensure that the dynamic crypto-map is enabled on this topology. In a full mesh VPN topology, you can apply only static crypto map policies.

IKEv2 Mode

For IPsec IKEv2 only, specify the encapsulation mode for applying ESP encryption and authentication to the tunnel. This determines what part of the original IP packet has ESP applied.

- **Tunnel mode**—(default) Encapsulation mode is set to tunnel mode. Tunnel mode applies ESP encryption and authentication to the entire original IP packet (IP header and data), hiding the ultimate source and destination addresses and becoming the payload in a new IP packet.

  The major advantage of tunnel mode is that you don't need to modify the end systems to receive the benefits of IPsec. This mode allows a network device, such as a router, to act as an IPsec proxy. That is, the router performs encryption on behalf of the hosts. The source router encrypts packets and forwards them along the IPsec tunnel. The destination router decrypts the original IP datagram and forwards it onto the destination system. Tunnel mode also protects against traffic analysis; with tunnel mode, an attacker can only determine the tunnel endpoints and not the true source and destination of the tunneled packets, even if they are the same as the tunnel endpoints.

- **Transport preferred**—Encapsulation mode is set to transport mode with an option to fall back to tunnel mode if the peer doesn't support it. In transport mode only the IP payload is encrypted, and the original IP headers are left intact. Therefore, the admin must select a protected network that matches the VPN interface IP address.

  This mode has the advantages of adding only a few bytes to each packet and allowing devices on the public network to see the final source and destination of the packet. With transport mode, you can enable special processing (for example, QoS) on the intermediate network based on the information in the IP header. However, the layer 4 header is encrypted, which limits examination of the packet.

- **Transport required**— Encapsulation mode is set to transport mode only, falling back to tunnel mode is allowed. If the endpoints can't successfully negotiate transport mode, due to one endpoint not supporting it, the VPN connection is not made.

Proposals

Click **Edit** (✎) to specify the proposals for your chosen IKEv1 or IKEv2 method. Select from the available **IKEv1 IPsec Proposals** or **IKEv2 IPsec Proposals** objects, or create and then select a new one. See

Configure IKEv1 IPsec Proposal Objects, on page 1473 and Configure IKEv2 IPsec Proposal Objects, on page 1474 for details.

Enable Security Association (SA) Strength Enforcement

Enabling this option ensures that the encryption algorithm used by the child IPsec SA isn't stronger (in terms of the number of bits in the key) than the parent IKE SA.

Enable Reverse Route Injection

Reverse Route Injection (RRI) enables static routes to be automatically inserted into the routing process for those networks and hosts protected by a remote tunnel endpoint.

Enable Perfect Forward Secrecy

Whether to use Perfect Forward Secrecy (PFS) to generate and use a unique session key for each encrypted exchange. The unique session key protects the exchange from subsequent decryption, even if the entire exchange was recorded and the attacker has obtained the preshared or private keys used by the endpoint devices. If you select this option, also select the Diffie-Hellman key derivation algorithm to use when generating the PFS session key in the Modulus Group list.

Modulus Group

The Diffie-Hellman group to use for deriving a shared secret between the two IPsec peers without transmitting it to each other. A larger modulus provides higher security but requires more processing time. The two peers must have a matching modulus group. For a full explanation of the options, see Deciding Which Diffie-Hellman Modulus Group to Use, on page 1503.

Lifetime Duration

The number of seconds a security association exists before expiring. The default is 28,800 seconds.

Lifetime Size

The volume of traffic (in kilobytes) that can pass between IPsec peers using a given security association before it expires. The default is 4,608,000 kilobytes. Infinite data isn't allowed.

ESPv3 Settings

Validate incoming ICMP error messages

Choose whether to validate ICMP error messages received through an IPsec tunnel and destined for an interior host on the private network.

Enable 'Do Not Fragment' Policy

Define how the IPsec subsystem handles large packets that have the do-not-fragment (DF) bit set in the IP header.

Policy

- Copy DF bit—Maintains the DF bit.
- Clear DF bit—Ignores the DF bit.
- Set DF bit—Sets and uses the DF bit.

Enable Traffic Flow Confidentiality (TFC) Packets

Enable dummy TFC packets that mask the traffic profile which traverses the tunnel. Use the **Burst**, **Payload Size**, and **Timeout** parameters to generate random length packets at random intervals across the specified SA.

Note You can enable dummy Traffic Flow Confidentiality (TFC) packets at random lengths and intervals on an IPsec security association (SA). You must have an IKEv2 IPsec proposal set before enabling TFC.

Enabling TFC packets prevents the VPN tunnel from being idle. Thus the VPN idle timeout configured in the group policy doesn't work as expected when you enable the TFC packets.

Firewall Threat Defense Advanced Site-to-site VPN Deployment Options

The following sections describe the advanced options you can specify in your site-to-site VPN deployment. These settings apply to the entire topology, all tunnels, and all managed devices.

Firewall Threat Defense VPN Advanced IKE Options

Advanced > IKE > ISAKAMP Settings

IKE Keepalive

Enable or disables IKE Keepalives. You can set this option to EnableInfinite so that the device never starts the keepalive monitoring itself.

Threshold
Specifies the IKE keep alive confidence interval. This interval is the number of seconds allowing a peer to idle before beginning keepalive monitoring. The minimum and default interval is 10 seconds; the maximum interval is 3600 seconds.

Retry Interval
Specifies number of seconds to wait between IKE keep alive retries. The default is 2 seconds, the maximum is 10 seconds.

Identity Sent to Peers:
Choose the identity that the peers will use to identify themselves during IKE negotiations:

- **autoOrDN**(default)—Determines IKE negotiation by connection type: IP address for preshared key, or Cert DN for certificate authentication (not supported).
- **ipAddress**—Uses the IP addresses of the hosts exchanging ISAKMP identity information.
- **hostname**—Uses the fully qualified domain name of the hosts exchanging ISAKMP identity information. This name comprises the hostname and the domain name.

Note Enable or disable this option for all your VPN connections.

Peer Identity Validation
During IKE tunnel establishment, the peer provides its identity: either an IP address, a Fully Qualified Domain Name (FQDN), or a Distinguished Name (DN). It also presents a certificate, which contains none, some, or all of these fields.

If IKE peer identity validation is enabled, the Firewall Threat Defense compares the peer's identity to the respective field in the certificate to see if the information matches. If the information matches, then

the peer's identity is validated and the Firewall Threat Defense establishes the tunnel. If the information does not match, the tunnel is not established.

- **Do not check**—Firewall Threat Defense does not validate the peer identity.
- **Required**—Firewall Threat Defense validates the peer identity.
- **If supported by cert**—Firewall Threat Defense validates the peer identity only if the peer provides a certificate.

Enable Aggressive Mode
Select this negotiation method for exchanging key information if the IP address isn't known and DNS resolution might not be available on the devices. Negotiation is based on hostname and domain name.

Enable Notification on Tunnel Disconnect
Allows an administrator to enable or disable the sending of an IKE notification to the peer when an inbound packet that is received on an SA does not match the traffic selectors for that SA. This notification is disabled by default.

Advanced > IKE > IVEv2 Security Association (SA) Settings

More session controls are available for IKE v2 that limit the number of open SAs. By default, there's no limit to the number of open SAs.

Cookie Challenge
Whether to send cookie challenges to peer devices in response to SA initiate packets, which can help thwart denial of service (DoS) attacks. The default is to use cookie challenges when 50% of the available SAs are in negotiation. Select one of these options:

- Custom
- Never (default)
- Always

Threshold to Challenge Incoming Cookies
The percentage of the total allowed SAs that are in-negotiation. This triggers cookie challenges for any future SA negotiations. The range is zero to 100%.

Number of SAs Allowed in Negotiation
Limits the maximum number of SAs that can be in negotiation at any time. If used with Cookie Challenge, configure the cookie challenge threshold lower than this limit for an effective cross-check.

Maximum number of SAs Allowed
Limits the number of allowed IKEv2 connections. Default is unlimited.

Enable Notification on Tunnel Disconnect
Allows an administrator to enable or disable the sending of an IKE notification to the peer when an inbound packet that is received on an SA doesn't match the traffic selectors for that SA. Sending this notification is disabled by default.

Firewall Threat Defense VPN Advanced IPsec Options

Advanced > IPsec > IPsec Settings

Enable Fragmentation Before Encryption
This option lets traffic travel across NAT devices that don't support IP fragmentation. It doesn't impede the operation of NAT devices that do support IP fragmentation.

Path Maximum Transmission Unit Aging
Check to enable Path Maximum Transmission Unit (PMTU) Aging, the interval to reset the PMTU of a Security Association (SA).

Value Reset Interval
Enter the number of minutes at which the PMTU value of an SA is reset to its original value. Valid range is 10 to 30 minutes, default is unlimited.

Firewall Threat Defense Advanced Site-to-site VPN Tunnel Options

Navigation Path

Advanced > Tunnel

Tunnel Options

Only available for Hub and Spoke, and Full Mesh topologies. This section doesn't appear for Point to Point configurations.

- **Enable Spoke to Spoke Connectivity through Hub**—Disabled by default. Choosing this field enables the devices on each end of the spokes to extend their connection through the hub node to the other device.

NAT Settings

- **Keepalive Messages Traversal** —Enabled by default. This parameter is a global setting that enables NAT-T for all endpoints within a topology. Check this check box to enable keepalive messages for NAT traversal. NAT traversal keepalive is used for the transmission of keepalive messages when there's a device (middle device) located between a VPN-connected hub and spoke, and that device performs NAT on the IPsec flow.

  NAT traversal allows seamless communication between the peer Firewall Threat Defense devices when there are NAT devices between these devices. For hub and spoke topologies, this option is available only for the spoke.

  If you select this option, configure the **Interval**, in seconds, between the keepalive signals sent between the spoke and the middle device to indicate that the session is active. The value can be from 5 to 3600 seconds. The default is 20 seconds. You can disable this feature for an endpoint in a topology when you add it using the VPN wizard (**Enable NAT Traversal** check box in the **Add Endpoint** dialog box).

Session Settings

- **Enable VPN Idle Timeout**—Enabled by default. This timeout is the amount of time that the VPN connection can be idle, without any activity in the tunnel, before it is disconnected. The default timeout is 30 minutes.

Access Control for VPN Traffic

Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)—By default, the Firewall Threat Defense applies access control policy inspection on the decrypted traffic. Enable this option to bypass the ACL inspection. The Firewall Threat Defense still applies the VPN Filter ACL and authorization ACL downloaded from the AAA server to the VPN traffic.

Enable or disable the option for all your VPN connections. If you disable this option, ensure that the traffic is allowed by the access control policy or prefilter policy.

Note For route-based VPNs, **sysopt permit-vpn** does not work. You must always create access control rules to allow route-based VPN traffic.

Certificate Map Settings

- **Use the certificate map configured in the Endpoints to determine the tunnel**—If this option is enabled (checked), the tunnel is determined by matching the contents of the received certificate to the certificate map objects configured in the endpoint nodes.
- **Use the certificate OU field to determine the tunnel**—Indicates that if a node isn't determined based on the configured mapping (the above option) if selected, then use the value of the organizational unit (OU) in the subject distinguished name (DN) of the received certificate to determine the tunnel.
- **Use the IKE identity to determine the tunnel**—Indicates that if a node isn't determined based on a rule matching or taken from the OU (the above options) if selected, then the certificate-based IKE sessions are mapped to a tunnel based on the content of the phase1 IKE ID.
- **Use the peer IP address to determine the tunnel**—Indicates that if a tunnel isn't determined based on a rule matching or taken from the OU or IKE ID methods (the above options) if selected, then use the established peer IP address.

About Virtual Tunnel Interfaces

Firewall Management Center supports a routable logical interface called the Virtual Tunnel Interface (VTI). VTIs do not require a static mapping of IPsec sessions to a physical interface. The IPsec tunnel endpoint is associated with a virtual interface. You can use these interfaces like other interfaces and apply static and dynamic routing policies.

As an alternative to policy-based VPN, you can create a VPN tunnel between peers using VTIs. VTIs support route-based VPN with IPsec profiles attached to the end of each tunnel. VTIs use static or dynamic routes. The device encrypts or decrypts the traffic from or to the tunnel interface and forwards it according to the routing table. Deployments become easier, and having VTI which supports route-based VPN with dynamic routing protocol also satisfies many requirements of a virtual private cloud. Firewall Management Center enables you to easily migrate from crypto-map based VPN configuration to VTI-based VPN.

You can configure route-based VPN with static or dynamic VTI using the site-to-site VPN wizard. Traffic is encrypted using static route, BGP, OSPFv2/v3, or EIGRP.

You can create a routed security zone, add VTI interfaces to it, and define access control rules for the decrypted traffic control over the VTI tunnel.

You can create VTI-based VPNs between:

- Two Firewall Threat Defense devices.
- A Firewall Threat Defense and public cloud.
- One Firewall Threat Defense and another Firewall Threat Defense with service provider redundancy.
- A Firewall Threat Defense and any other device with VTI interfaces.
- A Firewall Threat Defense and another device with policy-based VPN configuration.

There are two types of VTI interfaces: static VTI and dynamic VTI.

For more information, see Static VTI, on page 1530 and Dynamic VTI, on page 1531.

Static VTI

Static VTI uses tunnel interfaces to create a tunnel that is always-on between two sites. You must define a physical interface as a tunnel source for a static VTI. You can associate a maximum of 1024 VTIs per device. To create a static VTI interface in the management center, see Add a VTI Interface, on page 1537.

The figure below shows a VPN topology using static VTIs.

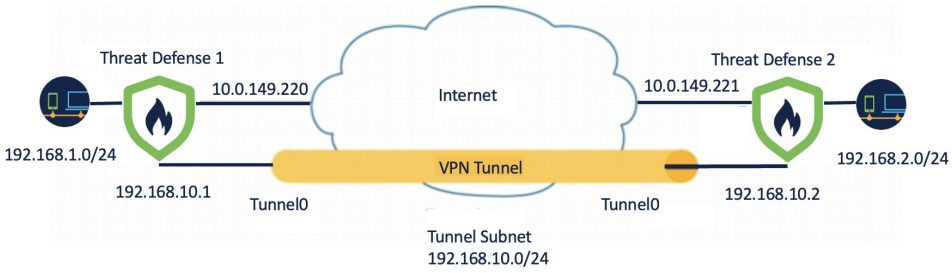

On Threat Defense 1:

- Static VTI IP address is 192.168.10.1
- Tunnel source is 10.0.149.220
- Tunnel destination is 10.0.149.221

On Threat Defense 2:

- Static VTI IP address is 192.168.10.2
- Tunnel source is 10.0.149.221
- Tunnel destination is 10.0.149.220

Benefits

- Minimizes and simplifies configuration.

  You do not have to track all remote subnets for a crypto map access list, and configure complex access lists or crypto maps.

- Provides a routable interface.

  Supports IP routing protocols such as BGP, EIGRP, and OSPFv2/v3, and static routes.

- Supports backup VPN tunnels
- Supports load balancing using ECMP.
- Supports virtual routers.
- Provides differential access control for VPN traffic.

  You can configure a VTI with a security zone and use it in an AC policy. This configuration:

- Allows you to classify and differentiate VPN traffic from clear-text traffic and permit VPN traffic selectively.
- Provides differential access-control for VPN traffic across different VPN tunnels.

Dynamic VTI

Dynamic VTI uses a virtual template for dynamic instantiation and management of IPsec interfaces. The virtual template dynamically generates a unique virtual access interface for each VPN session. Dynamic VTI supports multiple IPsec security associations and accepts multiple IPsec selectors proposed by the spoke.

Benefits

- Minimizes and simplifies configuration.

  You do not have to configure complex access lists or crypto maps.

- Simplifies management.
  - Easily manage peer configuration for large enterprise hub and spoke deployments.
  - Use only one dynamic VTI for multiple spokes, instead of configuring one static VTI per spoke.
- Provides a routable interface.

  Supports IP routing protocols such as BGP, EIGRP, and OSPFv2/v3, and static routes.

- Simplifies scaling

  Addition of new spokes does not require any additional VPN configuration on the hub. You may need to update NAT and routing configurations depending upon the setup.

- Support for backup VPN tunnels.
- Supports dynamic spokes.

  You do not have to update the hub configuration for spoke's DHCP IP address changes.

- Conserves IP addresses.
  - Uses the IP unnumbered interface functionality to borrow the IP address from another physical or loopback interface.
  - All virtual access interfaces associated to a dynamic VTI use the same IP address.
- Supports virtual routers.
- Provides differential access control for VPN traffic.

  You can configure a VTI with a security zone and use it in an AC policy. This configuration:
  - Allows you to classify and differentiate VPN traffic from clear-text traffic and permit VPN traffic selectively.
  - Provides differential access-control for VPN traffic across different VPN tunnels.

How Does the Firewall Management Center Create a Dynamic VTI Tunnel for a VPN Session

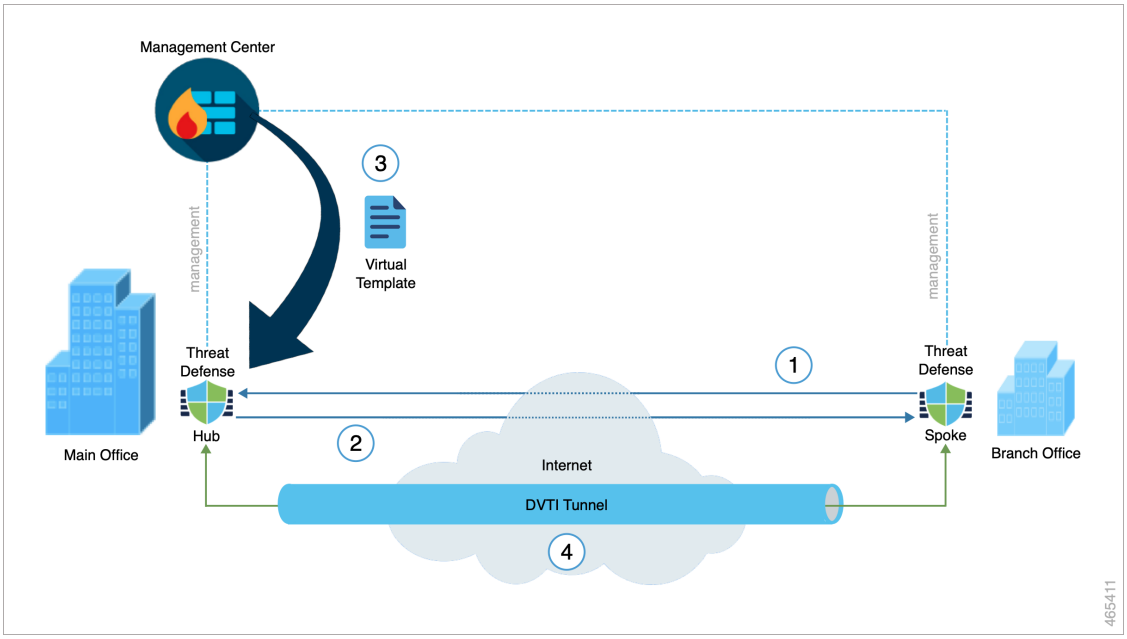

When a spoke initiates a tunnel request with the hub:

1. The spoke initiates an IKE exchange with the hub for a VPN connection.
2. The hub authenticates the spoke.
3. The Firewall Management Center assigns a dynamic virtual template on the hub for the spoke.

   The virtual template dynamically generates a virtual access interface on the hub. This interface is unique for the VPN session with the spoke.
4. The hub establishes a dynamic VTI tunnel with the spoke using the virtual access interface.
   a. The hub and spoke exchange traffic over the tunnel using:
      - Specific traffic proposed by the spokes over IKE exchanges.
      - BGP/OSPF/EIRGP protocols over the IPsec tunnel.
   b. After the VPN session ends, the tunnel disconnects and the hub deletes the corresponding virtual access interface.

To create a dynamic VTI interface in the Firewall Management Center, see Add a VTI Interface, on page 1537.

To configure a route-based site-to-site VPN using dynamic VTI, see Configure Dynamic VTI for a Route-based Site-to-Site VPN, on page 1551.

Virtual Routers and Dynamic VTI

You can create virtual routers, associate dynamic VTIs with these virtual routers, and extend the capabilities of dynamic VTIs in your network. You can associate dynamic VTIs either with global or user-defined virtual routers. You can assign a dynamic VTI to only one virtual router.

A virtual router associated with:

- A dynamic VTI is called an Indoor VRF (IVRF).
- A tunnel source interface is known as Front Door VRF (FVRF).

A dynamic VTI and its corresponding protected network interface must be part of the same virtual router. You must map the borrow IP interface and the dynamic VTI to the same virtual router. A tunnel source interface can be part of multiple virtual routers.

To configure virtual routers using dynamic VTI for a route-based site-to-site VPN, see How to Configure a Virtual Router with Dynamic VTI, on page 1168.

For more information about a configuration example, see How to Secure Traffic from Networks with Multiple Virtual Routers over a Site-to-Site VPN with Dynamic VTI, on page 1201

Guidelines and Limitations for Virtual Tunnel Interfaces

IPv6 Support

- VTI supports IPv6.
- You can use an IPv6 address for the tunnel source interface and use the same address as the tunnel endpoint.
- The Firewall Management Center supports the following combinations of VTI IP (or internal networks IP version) over public IP versions:
  - IPv6 over IPv6
  - IPv4 over IPv6
  - IPv4 over IPv4
  - IPv6 over IPv4
- VTI supports static and dynamic IPv6 addresses as the tunnel source and destination.
- The tunnel source interface can have IPv6 addresses and you can specify the tunnel endpoint address. If you don’t specify the address, by default, the Firewall Threat Defense uses the first IPv6 global address in the list as the tunnel endpoint.

BGP IPv6 Support

VTI supports IPv6 BGP.

EIGRP IPv4 Support

VTI supports IPv4 EIGRP.

OSPFv2 and OSPFv3 IPv6/IPv4 Support

VTI supports IPv4 and IPv6 OSPF.

ECMP Support

- Configure the spokes' static VTIs in an ECMP zone to load balance the application traffic. If you do not configure the ECMP zone, the remaining paths act as backup paths when the primary path goes down.

Multi-instance and Clustering

- VTI is supported in multi-instance.
- VTIs aren’t supported with clustering.

Firewall Mode

VTI is supported in routed mode only.

Limitations for Static VTI

- Only 20 unique IPSec profiles are supported.
- In route-based routing, you can configure VTI only as an egress interface.

Limitations for Dynamic VTI

- Dynamic VTI does not support:
  - ECMP
  - VRF in multi-instance
  - Clustering
  - IKEv1
  - QoS
- If a spoke has a dynamic IP address and a hub has a dynamic VTI behind a NAT, the tunnel status will be unknown.
- For a dynamic extranet, when multiple spokes establish a connection, the site-to-site monitoring dashboard does not show the individual tunnels.
- If you configure a hub with dynamic VTI behind NAT with dynamic spokes, the VPN monitoring data will not be accurate.

General Configuration Guidelines for Static and Dynamic VTI

- If you use dynamic crypto maps and dynamic VTIs in your site-to-site VPNs, only the dynamic VTI tunnels will come up. This behaviour occurs because both the crypto maps and dynamic VTIs try to use the default tunnel group.

  We recommend that you do one of the following:

  - Migrate your site-to-site VPNs to dynamic VTIs.
  - Use static crypto maps with their own tunnel-groups.

- VTIs are only configurable in IPsec mode.
- Dynamic VTI supports only the hub and spoke topology in the management center.
- Dynamic VTI supports only threat defense devices from version 7.3.
- We recommend that you configure only a single hub for a route-based hub and spoke topology. To configure a topology with multiple hubs for a set of spokes, with one hub as the backup hub, configure multiple topologies with a single hub and the same set of spokes. For more information, see Configure Multiple Hubs in a Route-based VPN, on page 1545.
- You can use static, BGP, EIGRP IPv4, OSPFv2/v3 routes for traffic using the tunnel interface.
- In an HA configuration with dynamic routing, the standby device cannot access the known subnets through the VTI tunnels as these tunnels are created with the active IP address.
- You can configure a maximum of 1024 static and dynamic VTIs on a device. While calculating the VTI count, consider the following:
  - Include nameif subinterfaces to derive the total number of VTIs that can be configured on the device.
  - You can't configure nameif on the member interfaces of a portchannel. Therefore, the tunnel count is reduced by the count of actual main portchannel interfaces alone and not any of its member interfaces.
  - The VTI count on a platform is limited to the number of VLANs configurable on that platform. For example, Firepower 1120 supports 512 VLANs, the tunnel count is 512 *minus* the number of physical interfaces configured.
- If you're configuring more than 400 VTIs on a device in a high-availability setup, you must configure 45 seconds as the unit holdtime for the Firewall Threat Defense HA.
- The MTU for VTIs is automatically set, according to the underlying physical interface.
- For dynamic VTI, the virtual access interface inherits the MTU from the configured tunnel source interface. If you don't specify the tunnel source interface, the virtual access interface inherits the MTU from the source interface from which the threat defense accepts the VPN session request.
- Static VTI supports IKE versions v1, v2, and uses IPsec for sending and receiving data between the tunnel's source and destination.
- Dynamic VTI supports only IKE version v2, and uses IPsec for sending and receiving data between the tunnel's source and destination.
- For static and dynamic VTI, ensure that you don't use the borrow IP interface as the tunnel source IP address for any VTI interface.
- When you configure a route-based site-to-site VPN using static or dynamic VTI interfaces, ensure that the value of the TTL hop is more than one if you use BGP.
- If NAT has to be applied, the IKE and ESP packets are encapsulated in the UDP header.
- IKE and IPsec security associations are re-keyed continuously regardless of data traffic in the tunnel. This ensures that VTI tunnels are always up.
- Tunnel group name must match what the peer sends as its IKEv1 or IKEv2 identity.

- For IKEv1 in LAN-to-LAN tunnel groups, you can use names which aren't IP addresses, if the tunnel authentication method is digital certificates and/or the peer is configured to use aggressive mode.
- VTI and crypto map configurations can coexist on the same physical interface, if the peer address configured in the crypto map and the tunnel destination for the VTI are different.
- By default, all traffic sent through a VTI is encrypted.
- Access rules can be applied on a VTI interface to control traffic through VTI.
- You can associate VTI interfaces with ECMP zones and configure ECMP static routes to achieve the following:
  - Load balancing (Active/Active VTIs)—Connection can flow over any of the parallel VTI tunnels.
  - Seamless connection migration—When a VTI tunnel becomes unreachable, the flows are seamlessly migrated to another VTI interface that is configured in the same zone.
  - Asymmetric routing—Forward traffic flow through one VTI interface and configure the reverse traffic flow through another VTI interface.

  For information on configuring ECMP, see Configure an Equal Cost Static Route, on page 1226.
- For route-based VPNs, Bypass Access Control policy for decrypted traffic (**sysopt connection permit-vpn**) does not work. You must always create access control rules to allow route-based VPN traffic.
- Ensure that you remove unused transform sets from route-based IKE VPN configurations. If you have unused transform sets, the VPN tunnels may not come up and encapsulation and decapsulation will not happen. Use `show run crypto` and `show vpn-sessiondb detail l2l` commands to view the transform sets.

Backup VTI Guidelines and Limitations

- Flow resiliency across tunnel failovers isn't supported. For example, the clear text TCP connection gets lost after a tunnel failover, and you need to reinitiate any FTP transfer that took place during the failover.
- Certificate authentication isn't supported in backup VTI.

Guidelines for Dynamic VTI and Virtual Routers

- A dynamic VTI and its corresponding protected network interface must be part of the same virtual router.
- You must map the borrow IP interface and the dynamic VTI to the same virtual router.
- User-defined virtual routers support only BGPv4/v6 and OSPFv2 routing protocols.
- A tunnel source interface can be in a different user-defined virtual router than that associated with the dynamic VTI.

Related Topics

Guidelines and Limitations for Loopback Interfaces, on page 829
Create a Route-based Site-to-Site VPN, on page 1538

Add a VTI Interface

For configuring a route-based site-to-site VPN, you must create a VTI interface on the devices at both the nodes of the VTI tunnel.

When you specify the tunnel type as dynamic and configure the related parameters, the Firewall Management Center generates a dynamic virtual template. The virtual template dynamically generates the virtual access interface that is unique for each VPN session.

Before you begin

Configure a loopback interface for redundancy of static and dynamic VTI VPN tunnels. For more information, see Configure a Loopback Interface, on page 830.

For a Secure Firewall 1200, Secure Firewall 3100 or Secure Firewall 4200 device, IPsec flow offload is also used when the device's VTI loopback interface is enabled.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click the **Edit** icon next to the device on which you want to create a VTI interface.

Step 3 Choose **Add Interfaces > Virtual Tunnel Interface**.

Step 4 Select the **Tunnel Type** as **Static** or **Dynamic**.

Step 5 Enter the name and description for the interface. By default, the interface is enabled.

Ensure that you specify a name that is not longer than 28 characters.

Step 6 (Optional) Choose a security zone from the **Security Zone** drop-down list to add the static VTI or dynamic VTI interface to that zone.

If you want to perform traffic inspection based on a security zone, add the VTI interface to the security zone and configure an access control (AC) rule. To permit the VPN traffic over the tunnel, you need to add an AC rule with this security zone as the source zone.

Step 7 In the **Priority** field, enter the priority to load balance the traffic across multiple VTIs.

The range is from 0 to 65535. The lowest number has the highest priority. This option is not applicable for dynamic VTI.

Step 8 Depending on the tunnel type, do one of the following:

- For a dynamic VTI, enter a unique ID in the range of 1 to 10413 in the **Template ID** field.
- For a static VTI, enter a unique tunnel ID in the range of 0 to 10413 in the **Tunnel ID** field.

Step 9 (Optional for dynamic VTI) Choose the tunnel source interface from the **Tunnel Source** drop-down list.

The VPN tunnel terminates at this interface, a physical or loopback interface. Choose the IP address of the interface from the drop-down list. You can select the IP address irrespective of the IPsec tunnel mode. In case of multiple IPv6 addresses, select the address that you want to use as the tunnel endpoint.

Step 10 Under **IPSec Tunnel Mode**, click the **IPv4** or **IPv6** radio button to specify the traffic type over the IPsec tunnel.

Step 11 Under **IP Address**:

- **Configure IP**: Enter the IPv4 or IPv6 address for the static VTI interface. You cannot configure an IP address for a dynamic VTI interface. Use the **Borrow IP** field for the dynamic VTI interface.
- **Borrow IP (IP unnumbered)**: Choose a physical or loopback interface from the drop-down list, the VTI interface inherits this IP address.

  Ensure that you use an IP address different from the tunnel source IP address. You can use this option for a static or dynamic VTI interface.

  Click Add + to configure a loopback interface. The loopback interface helps to overcome path failures. If an interface goes down, you can access all interfaces through the IP address assigned to the loopback interface.

Step 12 Click **OK**.

Step 13 Click **Save**.

Create a Route-based Site-to-Site VPN

You can configure a route-based site-to-site VPN for the following two topologies:

- **Point to Point** : Configure VTIs on both nodes of the tunnel and use the wizard to configure the VPN.
- **Hub and Spoke**: Configure VTIs on the hub and the spokes. Configure the hub with a dynamic VTI and spokes with static VTIs.

You can configure an extranet device as the hub and managed devices as spokes. You can configure multiple hubs and spokes, and also configure backup hubs and spokes.

- For extranet hubs and spokes, you can configure multiple IPs as backup.
- For managed spokes, you can configure a backup static VTI interface along with the primary VTI interface.

For more information on VTI, see About Virtual Tunnel Interfaces, on page 1529.

Note All references to VTI stands for static VTI and dynamic VTI, unless mentioned.

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site**, and click **Add**.

Step 2 Enter a name for the VPN topology in the **Topology Name** field.

Step 3 Choose **Route Based (VTI)** and do one of the following:

- Select **Point to Point** as the network topology. To configure endpoints for a route-based **Point to Point** topology, see Configure Endpoints for a Point to Point Topology, on page 1539.
- Select **Hub and Spoke** as the network topology. To configure endpoints for a route-based **Hub and Spoke** topology, see Configure Endpoints for a Hub and Spoke Topology, on page 1542.

Step 4 Click **Create**.

Step 5 (Optional) Specify the **IKE** options for the deployment as described in Firewall Threat Defense VPN IKE Options, on page 1521.

Step 6 (Optional) Specify the **IPsec** options for the deployment as described in Firewall Threat Defense VPN IPsec Options, on page 1523.

Step 7 (Optional) Specify the **Advanced** options for the deployment as described in Firewall Threat Defense Advanced Site-to-site VPN Deployment Options, on page 1526.

Step 8 Click **Save**.

What to do next

After you configure VTI interfaces and VTI tunnel on both the devices, you must configure:

- A routing policy to route the VTI traffic between the devices over the VTI tunnel. For more information, see Configure Routing and AC Policies for VTI, on page 1552.
- An access control rule to allow encrypted traffic. Choose **Policies > Access Control**.

Configure Endpoints for a Point to Point Topology

Configure the following parameters to configure endpoints for a route-based site-to-site VPN for the **Point to Point** topology nodes:

Before you begin

Configure the basic parameters for a point-to-point topology in a route-based VPN as described in Create a Route-based Site-to-Site VPN, on page 1538 and click the **Endpoints** tab.

Procedure

Step 1 Under **Node A**, in the **Device** drop-down menu, select the name of the registered device (Firewall Threat Defense) or extranet as the first endpoint of your VTI tunnel.

For an extranet peer, specify the following parameters:

a. Specify the name of the device.

b. Enter the primary IP address in the **Endpoint IP address** field. If you configure a backup VTI, add a comma and, specify the backup IP address.

c. Click **OK**.

After configuring the above parameters for the extranet hub, specify the pre-shared key for the extranet in the **IKE** tab.

Note
The AWS VPC has **AES-GCM-NULL-SHA-LATEST** as the default policy. If the remote peer connects to AWS VPC, select **AES-GCM-NULL-SHA-LATEST** from the **Policy** drop-down list to establish the VPN connection without changing the default value in AWS.

Step 2 For a registered device, you can specify the VTI interface for Node A from the **Virtual Tunnel Interface** drop-down list.

The selected tunnel interface is the source interface for Node A and the tunnel destination for Node B.

If you want to create a new interface on Node A, click the Add + icon and configure the fields as described in Add a VTI Interface, on page 1537.

If you want to edit the configuration of an existing VTI, select the VTI in the **Virtual Tunnel Interface** drop-down field and click **Edit VTI**.

Step 3 If your Node A device is behind a NAT device, check the **Tunnel Source IP is Private** check box. In the **Tunnel Source Public IP Address** field, enter the tunnel source public IP address.

Step 4 **Send Local Identity to Peers**—Select this option to send local identity information to the peer device. Select one of the following **Local Identity Configuration** from the list and configure the local identity:

- **IP address**—Use the IP address of the interface for the identity.
- **Auto**—Use the IP address for pre-shared key and Cert DN for certificate-based connections.
- **Email ID**—Specify the email ID to use for the identity. The email ID can be up to 127 characters.
- **Hostname**—Use the fully qualified hostname.
- **Key ID**—Specify the key-id to use for the identity. The key ID must be fewer than 65 characters.

The local identity is used to configure a unique identity per IKEv2 tunnel, instead of a global identity for all the tunnels. The unique identity allows Firewall Threat Defense to have multiple IPsec tunnels behind a NAT to connect to a Cisco Umbrella Secure Internet Gateway (SIG).

For information about configuring a unique tunnel ID on Umbrella, see **Cisco Umbrella SIG User Guide**.

Step 5 (Optional) Click **Add Backup VTI** to specify an extra VTI as the backup interface and configure the parameters.

Note
Ensure that both peers of the topology do not have the same tunnel source for the backup VTI. A device cannot have two VTIs with the same tunnel source and tunnel destination; hence, configure a unique tunnel source and tunnel destination combination.

Though the virtual tunnel interface is specified under Backup VTI, the routing configuration determines which tunnel to be used as primary or backup.

Step 6 Under **Additional Configuration**, do the following:

- To route traffic to the VTI, click **Routing Policy**. Firewall Management Center displays the **Devices** > **Routing** page.

  You can configure the Static, BGP, OSPF v2/v3, or EIGRP routing for the VPN traffic.

- To permit VPN traffic, click **AC Policy**. Firewall Management Center displays the access control policy page of the device. Proceed to add an allow/block rule specifying the security zone of the VTI. If you configure a backup VTI, ensure to include the backup tunnel to the same security zone as that of the primary VTI. No specific settings are required for the backup VTI in the AC policy page.

Step 7 Expand **Advance Settings** to configure additional configurations for the device. For more information, see Advanced Configurations for a Point to Point Topology in a Route-based VPN, on page 1541.

Step 8 Repeat the above procedure for Node B.

Step 9 Click **OK**.

What to do next

- (Optional) Specify the **IKE** options for the deployment as described in Firewall Threat Defense VPN IKE Options, on page 1521.
- (Optional) Specify the **IPsec** options for the deployment as described in Firewall Threat Defense VPN IPsec Options, on page 1523.
- (Optional) Specify the **Advanced** options for the deployment as described in Firewall Threat Defense Advanced Site-to-site VPN Deployment Options, on page 1526.
- Click **Save**.
- To route traffic to the VTI, choose **Devices** > **Device Management**, edit the threat defense device and click the **Routing** tab.

  You can configure the static routes, or use BGP, OSPF v2/v3, or EIGRP for routing the VPN traffic.
- To permit VPN traffic, choose **Policies > Access Control.**. Add a rule specifying the security zone of the VTI. For a backup VTI, ensure that you include the backup VTI in the same security zone as that of the primary VTI.

Advanced Configurations for a Point to Point Topology in a Route-based VPN

Configure the following advanced configurations for a point-to-point topology in a route-based VPN:

Before you begin

Configure the basic parameters for a point-to-point topology in a route-based VPN as described in Configure Endpoints for a Point to Point Topology, on page 1539 and expand **Advance Settings**.

Procedure

Step 1 Check the **Send Virtual Tunnel Interface IP to the peers** check box to send the VTI IP address to the peer device.

Step 2 Check the **Allow incoming IKEv2 routes from the peers** check box to allow incoming IKEv2 routes from the spokes and peers.

Step 3 Choose one of the following from the **Connection Type** drop-down list:

Answer Only: The device can only respond when a peer device initiates a connection, it can't initiate any connection.

Bidirectional: The device can initiate or respond to a connection. This is the default option.

Configure Endpoints for a Hub and Spoke Topology

You can create a route-based site-to-site VPN using dynamic VTI only for hub and spoke topologies. The hub can use only a dynamic VTI and the spokes can use only static VTI interfaces. You can also configure an extranet device as a hub.

Configure the following parameters to configure endpoints for a route-based site-to-site VPN for the **Hub and Spoke** topology nodes:

Before you begin

Configure the basic parameters for a hub and spoke topology in a route-based VPN as described in Create a Route-based Site-to-Site VPN, on page 1538 and click the **Endpoints** tab.

Procedure

Step 1 Under **Hub Nodes**:

a) Click Add + to configure the hub node in the **Add Endpoint** dialog box.

b) Choose a hub from the **Device** drop-down list.

For an extranet hub, specify the following parameters:

1. Enter the name of the device.
2. Enter the primary IP address. If you configure a backup VTI, add a comma, and then specify the backup IP address.
3. Click **OK**.

After configuring the above parameters for the extranet hub, specify the pre-shared key for the extranet in the **IKE** tab.

Note
The AWS VPC has **AES-GCM-NULL-SHA-LATEST** as the default policy. If the remote peer connects to AWS VPC, select **AES-GCM-NULL-SHA-LATEST** from the **Policy** drop-down list to establish the VPN connection without changing the default value in AWS.

c) Choose a dynamic VTI from the **Dynamic Virtual Tunnel Interface** drop-down list.

Tunnel source configuration is mandatory for a dynamic VTI as the management center needs this information to determine the tunnel destination of the spoke.

Click Add + to add a new dynamic VTI. We recommend that you configure the Borrow IP for the dynamic interface from a loopback interface.

If you want to edit an existing dynamic VTI, select the interface, and click **Edit VTI**.

d) (Optional) If your endpoint device is behind a NAT device, check the **Tunnel Source IP is Private** check box and configure the tunnel source IP address in the **Tunnel Source Public IP Address** field.
e) Click **Routing Policy** to configure the routing policy for the hub.
f) Click **AC Policy** to configure the access control policy.
g) Expand **Advance Settings** to configure additional configurations on the hub. For more information, see Advanced Configurations for Hub and Spokes in a Route-based VPN, on page 1544.
h) Click **OK**.

Step 2 Under **Spoke Nodes**:

a) Click Add + to configure the spoke in the **Add Endpoint** dialog box.
b) Choose a spoke from the **Device** drop-down list.

For an extranet spoke, specify the following parameters:

1. Enter the name of the device.
2. Under **Endpoint IP Address**, choose one of the following:
   - **Static**: Enter the IP address of the device, and the backup IP address, if required.
   - **Dynamic**: Choose this option to dynamically assign the IP addresses for the extranet spokes.
3. Click **OK**.

c) Choose a static VTI from the **Static Virtual Tunnel Interface** drop-down list.

Click Add + to add a new static VTI. Tunnel IP of the static VTI is auto populated, ensure that this IP address is unique for the spoke.

If you want to edit an existing static VTI, select the interface, and click **Edit VTI**.

d) (Optional) If your endpoint device is behind a NAT device, check the **Tunnel Source IP is Private** check box. The management center needs the tunnel source interface address to configure the tunnel destination IP address on the spokes. In the **Tunnel Source Public IP Address** field, enter the tunnel source public IP address.
e) (Optional) **Send Local Identity to Peers**—Check this check box to send the local identity information to the peer device. Choose one of the following parameters from the **Local Identity Configuration** drop-down list and configure the local identity:
   - **IP address**—Use the IP address of the interface for the identity.
   - **Auto**—Use the IP address for pre-shared key and Cert DN for certificate-based connections.
   - **Email ID**—Specify the email ID to use for the identity. The email ID can be up to 127 characters.
   - **Hostname**—Use the fully qualified hostname.
   - **Key ID**—Specify the key-id to use for the identity. The key ID must be less than 65 characters.

The local identity is used to configure a unique identity per IKEv2 tunnel, instead of a global identity for all the tunnels. The unique identity allows Firewall Threat Defense to have multiple IPsec tunnels behind a NAT to connect to the Cisco Umbrella Secure internet Gateway (SIG).

For more information about configuring a unique tunnel ID on Umbrella, see *Cisco Umbrella SIG User Guide*.

f) (Optional) Click **Add Backup VTI** to specify an extra VTI interface as the backup interface.

Note
Ensure that both peers of the topology do not have backup VTI configured on the same tunnel source. For instance, if Peer A is having two VTIs (primary and a backup) configured with a single tunnel source interface, say, 10.10.10.1/30, then Peer B also can't have its 2 VTIs with a single tunnel source IP, say 20.20.20.1/30.

Though the virtual tunnel interface is specified under backup VTI, the routing configuration determines which tunnel to be used as primary or backup.

g) Click **Routing Policy** to configure the routing policy for the spoke.
h) Click **AC Policy** to configure the access control policy.
i) Expand **Advance Settings** to configure additional configurations on the spoke. For more information, see Advanced Configurations for Hub and Spokes in a Route-based VPN, on page 1544.
j) Click **OK**.

What to do next

- (Optional) Specify the **IKE** options for the deployment as described in Firewall Threat Defense VPN IKE Options, on page 1521.
- (Optional) Specify the **IPsec** options for the deployment as described in Firewall Threat Defense VPN IPsec Options, on page 1523.
- (Optional) Specify the **Advanced** options for the deployment as described in Firewall Threat Defense Advanced Site-to-site VPN Deployment Options, on page 1526.
- Click **Save**.

Advanced Configurations for Hub and Spokes in a Route-based VPN

Configure the following advanced configurations for a hub and spoke in a route-based VPN:

Before you begin

Configure the basic parameters for a hub and spoke in a route-based VPN as described in Configure Endpoints for a Hub and Spoke Topology, on page 1542 and expand **Advance Settings**.

Procedure

Step 1 Check the **Send Virtual Tunnel Interface IP to the peers** check box to send the VTI IP address to the peer device.

For a hub, you must check this check box if you use BGP as the routing protocol. This configuration ensures that the loopback IP address is shared in the BGP routing table.

For a spoke, this option is enabled by default.

Step 2 Add the **Protected Networks** to define the networks protected by the VPN endpoint. Click Add + to select a protected network.

For a hub, configure the protected networks behind the hub. This information and the spoke's protected network generate the spoke access list.

You cannot create a static route for a virtual access interface on a hub with dynamic VTI. The hub creates and deletes these interfaces dynamically during tunnel establishment and termination.

For a spoke, configure the spoke's protected network.

To enable static routing for the spokes, after you configure the endpoints for your topology, click the **IPsec** tab and check the **Enable Reverse Route Injection** check box.

You do not need this option if you use BGP, OSPF, or EIGRP.

Step 3 Check the **Allow incoming IKEv2 routes from the peers** check box to allow incoming IKEv2 routes from the spokes and peers.

For a hub: During an IKE exchange, the hub advertises the dynamically created virtual access interfaces to the spokes, and the spokes advertise their VTI IP addresses to the hub.

For a spoke: By default, this option is enabled.

Step 4 In the **Connection Type** drop-down list, choose one of the following:

Answer Only: The device can only respond when a peer device initiates a connection, it can’t initiate any connection.

Bidirectional: The device can initiate or respond to a connection. This is the default option.

Configure Multiple Hubs in a Route-based VPN

You can configure a topology with multiple hubs for a set of spokes. With one hub as the backup hub, you can configure multiple topologies with a single hub and the same set of spokes.

In the following example, there are two hubs connected to the same set of spokes. Hub 1 is the primary hub and Hub 2 is the secondary hub. To configure this network in the Firewall Management Center, you must configure two route-based hub and spoke topologies:

- Topology 1: Hub 1 connected to spoke 1 and spoke 2.
- Topology 2: Hub 2 connected to spoke 1 and spoke 2.

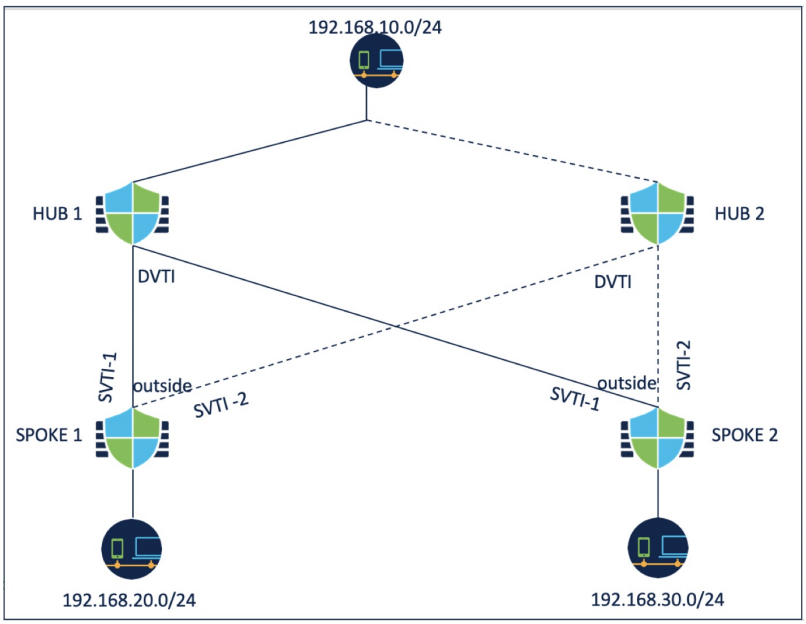

To configure topology 1:

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site**, and click **Add**.

Step 2 Enter a name for the VPN topology in the **Topology Name** field.

Step 3 Choose **Route Based (VTI)** and do one of the following:

- Select **Point to Point** as the network topology. To configure endpoints for a route-based **Point to Point** topology, see Configure Endpoints for a Point to Point Topology, on page 1539.
- Select **Hub and Spoke** as the network topology. To configure endpoints for a route-based **Hub and Spoke** topology, see Configure Endpoints for a Hub and Spoke Topology, on page 1542.

Step 4 Click **Create**.

Step 5 Configure the IKE version.

Step 6 Click the **Endpoints** tab.

Step 7 Under **Hub Nodes**:

a) Click Add + to add the hub.

b) Choose Hub 1 from the **Device** drop-down list.

c) Choose a dynamic VTI from the **Dynamic Virtual Tunnel Interface** drop-down list or click Add + to add a new dynamic VTI.

We recommend that you configure the Borrow IP for the dynamic interface from a loopback interface.

d) (Optional) If your endpoint device is behind a NAT device, check the **Tunnel Source IP is Private** check box and configure the tunnel source IP address in the **Tunnel Source Public IP Address** field.

e) Click **Routing Policy** to configure the routing policy for the hub. You can configure dynamic routing using BGP.

f) Expand **Advance Settings**. You can configure the following advanced settings for the hub to enable IKEv2 routing, which can be used if you do not use dynamic routing.

- (Optional) Check the **Send Virtual Tunnel Interface IP to the peers** check box.
- Check the **Allow incoming IKEv2 routes from the peers** check box for the hub to accept routes from the spokes and update the routing table.
- Choose **Connection Type** as Bidirectional from the drop-down list.

g) Click **OK**.

Step 8 Under **Spoke Nodes**:

a) Click Add + to add a spoke.
b) Choose Spoke 1 from the **Device** drop-down list.
c) Choose SVTI-1 as the static VTI for the spoke from the **Static Virtual Tunnel Interface** drop-down list or click Add + to add a new static VTI.

Choose the outside interface as the tunnel source of SVTI-1. Tunnel IP of SVTI-1 is autopopulated, ensure that this IP address is unique for spoke 1 across peers in both the topologies.

d) Expand **Advance Settings**. If you do not use dynamic routing, you can configure these settings to enable IKEv2 routing for the spoke.

- Check the **Send Virtual Tunnel Interface IP to the peers** check box to send the VTI IP address to the peer device.
- Check the **Allow incoming IKEv2 routes from the peers** check box to allow incoming IKEv2 routes from the peers.
- Choose **Connection Type** as Bidirectional from the drop-down list.

e) Click **OK**.
f) Repeat steps 5a to 5e to add spoke 2. Configure SVTI-1 as the static VTI of spoke 2.

Step 9 Configure the IKE and IPSec parameters as required or use the default values.

What to do next

1. Configure topology 2 with hub 2, spoke 1, and spoke 2.

   Configure SVTI-2 as the static VTI of spoke 1 and SVTI-2 as the static VTI of spoke 2 (refer the above illustration). Tunnel source for SVTI-2 should be the same outside interface.

2. For each spoke, configure the routing policy. For more information, see Configure Routing for Multiple Hubs in a Route-based VPN, on page 1547.

3. Verify the configuration and tunnel statuses. For more information, see Verify the Multiple Hubs Configuration in a Route-based VPN, on page 1549.

Configure Routing for Multiple Hubs in a Route-based VPN

The following procedure explains how to configure dynamic routing on the hub and spokes, and configure Policy Based Routing on the spokes.

Before you begin

Configure topology 1 and 2 as explained in Configure Multiple Hubs in a Route-based VPN, on page 1545.

Procedure

Step 1 Configure dynamic routing for the hub using BGP.

a) Choose **Device > Device Management > Routing**.
b) On the left pane, choose **General Settings > BGP**.
c) Check the **Enable BGP** check box and enter the **AS number**.

You can configure the other fields as per your requirement.

d) Click **Save**.
e) On the left pane, choose **BGP > IPv4**.
f) Check the **Enable IPv4** check box.
g) Click the **Neighbor** tab, click **Add** and configure the parameters.

1. **IP Address**: Enter the tunnel interface IP address of Spoke 1.
2. **Remote AS**: AS number of Spoke 1.
3. Check the **Enabled Address** check box.
4. Click **OK**.

Repeat the above steps to add Spoke 2 as a neighbor.

h) Click **Save**.
i) Click the **Networks** tab and click **Add** to advertise the network behind the hub to the peers.

Step 2 Configure dynamic routing for the spokes using BGP.

The BGP configuration for the spokes is similar to that of the hub except for the following differences:

- Configure Hub 1 and Hub 2 as the neighbors for both the spokes and use the tunnel interface IP address of the hubs.
- When you configure networks, use the network behind each spoke.

Step 3 Configure Policy Based Routing on the spokes.

a) On the left pane, choose **Policy Based Routing** and click **Add**.
b) Choose the **Ingress Interface** from the drop-down list.
c) Click **Add** to configure a Match ACL.

For example, for spoke 1, source network is 192.168.20.0/24 and destination network is 192.168.10.0/24.

d) Choose Egress Interfaces from the **Send to** drop-down list.
e) Choose Order from the **Interface Ordering** drop-down list.
f) Select the SVTI-1 and SVTI-2 interfaces as the egress interfaces.
g) Click **Save**.

If you want to use the hubs as a load-balancing pair, you must configure ECMP.

Step 4 Deploy the configurations on the hub and spokes.

What to do next

Verify the configurations and tunnel statuses. For more information, see Verify the Multiple Hubs Configuration in a Route-based VPN, on page 1549.

Verify the Multiple Hubs Configuration in a Route-based VPN

To verify the multiple hub configurations and the tunnel statuses:

- After the deployment, check the tunnel status in the dashboards.
- Use Packet Tracer from the Site-to-site monitoring dashboard to verify the selected path of the traffic (hub 1 or hub 2).
- Use the following show commands for each endpoint to verify the configurations:
  - **show run route-map**
  - **show run access-list**
  - **show route-map**
  - **show route**

Route Traffic Through a Backup VTI Tunnel

Secure Firewall Threat Defense supports the configuration of a backup tunnel for the route-based (VTI) VPN. When the primary VTI is unable to route the traffic, the traffic in the VPN is tunneled through the backup VTI.

You can deploy the backup VTI tunnel in the following scenarios:

- Both peers having service provider redundancy backup.

  In this case, there are two physical interfaces, acting as the tunnel sources for the two VTIs of the peers.
- Only one of the peers having service provider redundancy backup.

  In this case, there's an interface backup on only one side of the peer and on the other end, there is only one tunnel source interface.

Step	Do This	More Info
1	Review the guidelines and limitations.	Guidelines and Limitations for Virtual Tunnel Interfaces, on page 1533
2	Create the VTI interface.	Add a VTI Interface, on page 1537

Step	Do This	More Info
3	In the **Add Endpoint** dialog box of the **Create New VPN Topology** wizard, click **Add Backup VTI** to configure the respective backup interface for each peer.	• Configure Endpoints for a Point to Point Topology, on page 1539 • Configure Endpoints for a Hub and Spoke Topology, on page 1542
4	Configure the routing policy.	• Choose **Devices > Device Management**, and edit the threat defense device. • Click **Routing**.
5	Configure the access control policy.	• Choose **Policies > Access Control**.

Guidelines for Configuring a Backup VTI Tunnel

- For an extranet peer, you can specify the tunnel source IP address of the backup interface and configure the tunnel destination IP on the managed peer.

  You can specify the backup peer IP address in the **Endpoint IP Address** field of the **Create New VPN Topology** wizard.

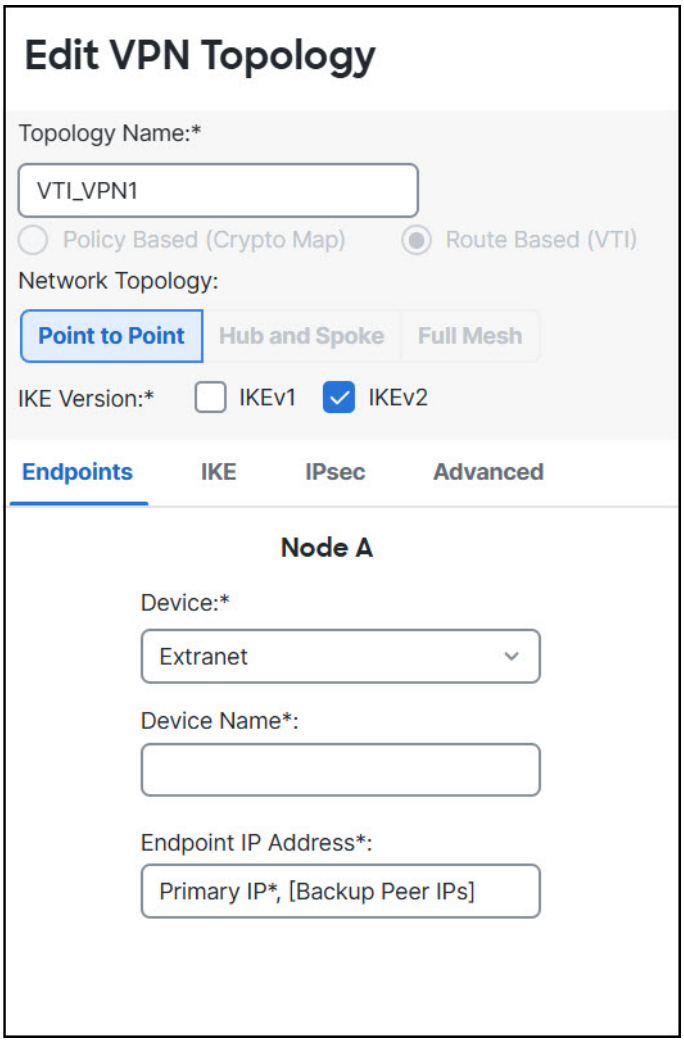

- After you configure the backup interfaces, configure the routing policy and access control policy for routing traffic.

  Though primary and backup VTIs are always available, traffic flows only through the tunnel that is configured in the routing policy. For detailed information, see Configure Routing and AC Policies for VTI, on page 1552.

- When you configure a backup VTI, ensure that you include the backup tunnel to the same security zone as that of the primary VTI. No specific settings are required for the backup VTI in the AC policy page.
- If you configure a static route for the backup tunnel, configure a static route with a different metric to handle the failover of the traffic flow over the backup tunnel.

Configure Dynamic VTI for a Route-based Site-to-Site VPN

To configure dynamic VTI for a route-based site-to-site VPN in the management center:

Step	Do This	More Information
1	Create a dynamic VTI interface on the hub.	Add a VTI Interface, on page 1537
2	Create static VTI interfaces on the spokes.	Add a VTI Interface, on page 1537
3	Create a route-based site-to-site VPN.	Create a Route-based Site-to-Site VPN, on page 1538
4	Configure the routing policy and the access control policy.	Configure Endpoints for a Hub and Spoke Topology, on page 1542

How to Configure a Virtual Router with Dynamic VTI

To configure a virtual router with dynamic VTI for a route-based site-to-site VPN in the management center:

Step	Do This	More Information
1	Create a route-based site-to-site VPN with a dynamic VTI interface on the hub and static VTI(s) on the spoke(s).	Create a Route-based Site-to-Site VPN, on page 1538
2	Create a virtual router.	Create a Virtual Router, on page 1178
3	Assign the interfaces to the virtual router.	Configure a Virtual Router, on page 1179
4	Configure routing policies for the hub and spoke.	Configure Endpoints for a Hub and Spoke Topology, on page 1542
5	Configure access control policies for the hub and spoke.	Configure Endpoints for a Hub and Spoke Topology, on page 1542

Configure Routing and AC Policies for VTI

After you configure VTI interfaces and the VTI tunnel on both the devices, you must configure:

- A routing policy to route VTI traffic between the devices over the VTI tunnel.
- An access control rule to allow encrypted traffic.

Routing Configuration for VTI

For the VTI interfaces, you can configure static route or routing protocols such as BGP, EIGRP, OSPF/OSPFv3.

1. Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
2. Click **Routing**.
3. Configure static route, or BGP, EIGRP, OSPF/OSPFv3.

Routing	Parameters	More Information
Static Route	• **Interface**—Select the VTI interface. For a backup tunnel, select the backup VTI interface. • **Selected Network**—Remote peer's protected network. • **Gateway**—Remote peer's tunnel interface IP address. For a backup tunnel, select the remote peer's backup tunnel interface IP address. • **Metric**—For a backup tunnel, configure a different metric to handle the failover of the traffic flow over the backup tunnel.	Add a Static Route, on page 1154

BGP	• Under **General Settings** > **BGP**, enable BGP, provide the AS number of the local device, and add Router ID (if you choose Manual).
	• Under **BGP**, enable IPv4/IPv6 and click the **Neighbor** tab to configure the neighbors.
	  • **IP Address**—Remote peer's VTI interface IP address. For a backup tunnel, add a neighbor with the remote peer's backup VTI interface IP address.
	  • **Remote AS**—Remote peer's AS number.
	• Click the **Redistribution** tab, select the **Source Protocol** as Connected to enable connected route redistribution.
		Configure BGP, on page 1281

Routing	Parameters	More Information
EIGRP	• Enable EIGRP, provide the AS number of the local device, and select the networks or hosts that participate in the EIGRP routing process. • Click the **Neighbors** tab and define the static neighbors for the EIGRP process. • To advertise summary addresses from a VTI interface, click the **Summary Address** tab, choose the VTI interface from the **Interface** drop-down. From the **Network** drop-down, choose the network to be summarized. • Click the **Interfaces** tab to configure the interface-specific EIGRP routing properties for the VTI interface. To enable EIGRP split-horizon on the interface, check the **Split Horizon** check box. You can also configure the **Hold Time** that is advertised by the device in the EIGRP hello packets.	Configure EIGRP, on page 1270
OSPF	• Check the **Process 1** check box, and choose the OSPF role. • Click the **Interface** tab and choose a VTI interface.	Configure OSPFv2, on page 1245
OSPFv3	• Check the **Process 1** and **Enable Process 1** check boxes, and choose the OSPFv3 role. • Click the **Interface** tab and choose a VTI interface.	Configure OSPFv3, on page 1257

AC Policy Rule

Add an access control rule to the access control policy on the device to allow encrypted traffic between the VTI tunnels with the following settings:

1. Create the rule with the Allow action.

2. Select the VTI security zone of the local device as the source zone and the VTI security zone of the remote peer as the destination zone.
3. Select the VTI security zone of the remote peer as the source zone and the VTI security zone of the local device as the destination zone.

For more information about configuring an access control rule, see Create and Edit Access Control Rules, on page 1783.

View Virtual Tunnel Information

You can view details of dynamic and static VTIs of route-based VPNs on the device. For every VPN topology, you can also view details of all the dynamically generated virtual access interfaces associated with each dynamic VTI.

Before you begin

- For static VTIs: Firewall Threat Defense Version 7.0 and later
- For dynamic VTIs: Firewall Threat Defense Version 7.3 and later

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** (✎) for your Firewall Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click the **Virtual Tunnels** tab.

For each VTI, you can view details such as name, IP address, IPsec mode, tunnel source interface details, topology, and remote peer IP. You can also view if path monitoring is enabled for each interface.

Deploy a SASE Tunnel on Umbrella

Cisco Umbrella is Cisco's cloud-based Secure Internet Gateway (SIG) platform that provides multiple levels of defense against internet-based threats. Umbrella integrates secure web gateway, DNS-layer security, and cloud access security broker (CASB) functionality to protect your systems against threats.

You can establish an IPsec IKEv2 tunnel from a threat defense device to Umbrella using the management center. This tunnel forwards all internet-bound traffic to the Umbrella SIG for inspection and filtering. This solution provides centralized security management so that network administrators don't have to separately manage the security settings of each branch.

To directly configure and deploy Umbrella tunnels from a threat defense device, you can create a SASE topology using a simple wizard. SASE topology is a new type of site-to-site VPN topology that supports:

- Static VTI-based site-to-site VPN.

- Hub and spoke topology, where Umbrella is the hub and the managed threat defense devices are the spokes.
- Pre-shared key based authentication.
- Firewall Threat Defense deployed in HA mode.
- Multi-instance: In a multi-instance deployment, you can integrate only one Umbrella account.

For high availability, you can configure two tunnels from a threat defense device and use the second tunnel as the backup tunnel. Ensure that you configure different local tunnel IDs for each tunnel.

For ease of configuration, the management center configures the default IPsec and IKEv2 policies.

Default IKEv2 policy configuration:

- Integrity Algorithm: NULL
- Encryption Algorithm: AES-GCM-256
- PRF Algorithm: SHA-256
- DH Group: 19, 20

Default IKEv2 IPsec policy configuration:

- ESP Hash: SHA-256
- ESP Encryption: AES-GCM-256

Related Topics

How to Deploy a SASE Tunnel on Umbrella, on page 1557

Guidelines and Limitations for Configuring SASE Tunnels on Umbrella

SASE topology supports:

- Only PSK-based authentication
- IKEv2
- High availability

General Configuration Guidelines

- The Firewall Management Center does not discover tunnels created directly on Umbrella or by other applications.
- You can add only devices managed by the Firewall Management Center as endpoints for the SASE topology. You cannot add extranet devices.

  For high availability pairs, the HA pair names appear in the endpoint list.

- When you delete a tunnel from the Firewall Management Center and if it is unable to delete the tunnel from Umbrella, you must manually delete it by logging into Umbrella.
- You cannot edit or delete a SASE topology if the deployment to Umbrella is in progress. You can view the tunnel deployment status in the:
  - Cisco Umbrella Configuration dialog box of the wizard
  - Notifications page under the Deployments and Tasks tabs
  - Site to Site VPN Monitoring dashboard
- If you check the **Deploy configuration on threat defense nodes** check box in the wizard, the Umbrella SASE topology configuration is deployed on the Firewall Threat Defense only after the tunnels are deployed on Umbrella.

  The Firewall Management Center requires the local tunnel ID to deploy the Umbrella configuration on the Firewall Threat Defense. Umbrella generates the complete tunnel ID (*<prefix>@<umbrella generated ID>*-umbrella.com) only after the Firewall Management Center deploys the tunnel on Umbrella.
- The Firewall Management Center does not recognize topologies with the Umbrella data center as an extranet hub, created before version 7.3, as SASE topologies. You must create new SASE topologies in version 7.3 and delete the existing topology.
- After a Firewall Threat Defense HA switchover, the SASE topology will not appear in the site-to-site monitoring/VPN summary dashboard. We recommend that you bring the tunnel down using the **vpn-sessiondb logoff index** command and bring it up using packet tracer.

Limitations

SASE topology does not support:

- Clustering
- Certificate-based authentication
- IKEv1

How to Deploy a SASE Tunnel on Umbrella

This section provides instructions to deploy a SASE tunnel on Umbrella from a Firewall Threat Defense device using the Firewall Management Center.

Step	Do This	More Info
1	Review the guidelines and limitations.	Guidelines and Limitations for Configuring SASE Tunnels on Umbrella, on page 1556
2	Ensure that you meet the prerequisites.	Prerequisites for Configuring Umbrella SASE Tunnels, on page 1558

Step	Do This	More Info
3	Configure Cisco Umbrella connection settings.	• Configure Cisco Umbrella Connection Settings • Map Management Center Umbrella Parameters and Cisco Umbrella API Keys, on page 1558
4	Configure a SASE tunnel on Umbrella.	Configure a SASE Tunnel for Umbrella, on page 1560
5	View the status of the SASE tunnel.	View SASE Tunnel Status, on page 1561

Prerequisites for Configuring Umbrella SASE Tunnels

- You must have a Cisco Umbrella Secure Internet Gateway (SIG) Essentials subscription.
- You must enable your Smart License account with the export-controlled features to deploy tunnels on Umbrella from the Firewall Management Center. If this license is not enabled, you can only create a SASE topology. You cannot deploy tunnels on Umbrella.
- You must establish an account with Cisco Umbrella at https://umbrella.cisco.com, log into Umbrella at http://login.umbrella.com, and obtain the required information to establish connection to Cisco Umbrella.
- You must register Cisco Umbrella with the Firewall Management Center and configure the management key and the management secret in the Cisco Umbrella Connection Settings. The management center requires the management key and management secret to fetch the datacenter details from the Cisco Umbrella cloud. You must also configure the Organization ID, Network Device Key, Network Device Secret, and the Legacy Network Device Token in the Cisco Umbrella Connection Settings.

  For more information, see:

  - Configure Cisco Umbrella Connection Settings
  - Map Management Center Umbrella Parameters and Cisco Umbrella API Keys, on page 1558
- Ensure that Umbrella data center is reachable from the Firewall Threat Defense.
- You can deploy a tunnel only between Cisco Umbrella and a Firewall Threat Defense version 7.1.0 and later.

Map Management Center Umbrella Parameters and Cisco Umbrella API Keys

To register the Cisco Umbrella with the Firewall Management Center and configure the Umbrella parameters in the Firewall Management Center, you must:

1. Log in to Cisco Umbrella.
2. Choose **Admin > API Keys > Legacy Keys**.
3. Generate and copy the required API keys.
4. Use the API keys to configure the Cisco Umbrella connection parameters in the Firewall Management Center.

The figure below shows the parameters that you must configure in the Cisco Umbrella Connection in the Firewall Management Center. DNScrypt Public Key is an optional parameter.

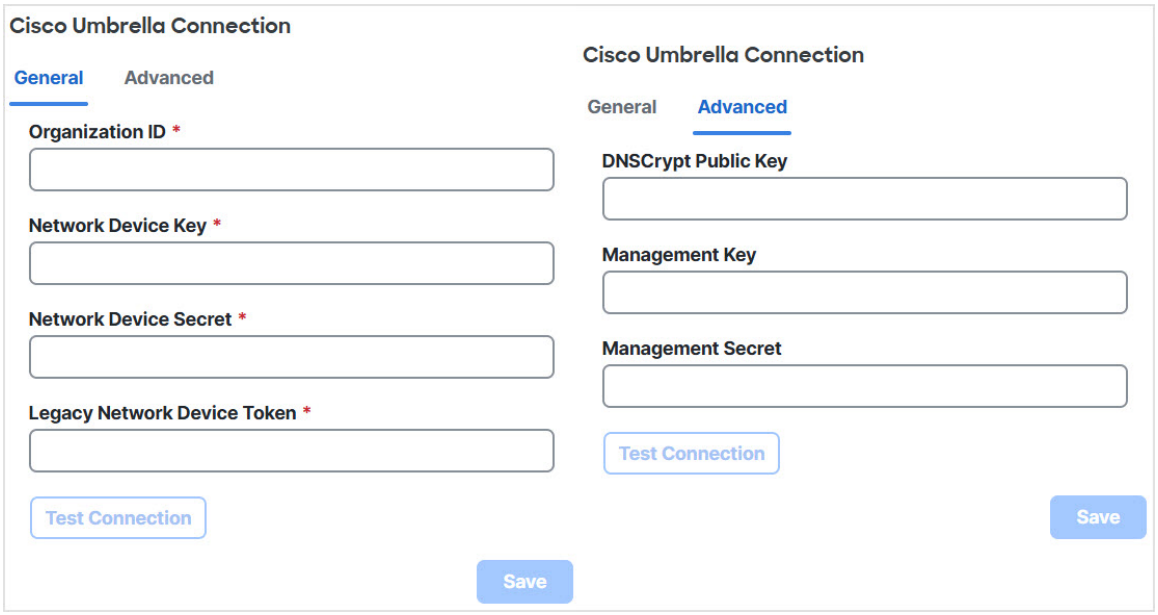

The figure below shows the Cisco Umbrella API keys that you must use to register the Cisco Umbrella with the Firewall Management Center.

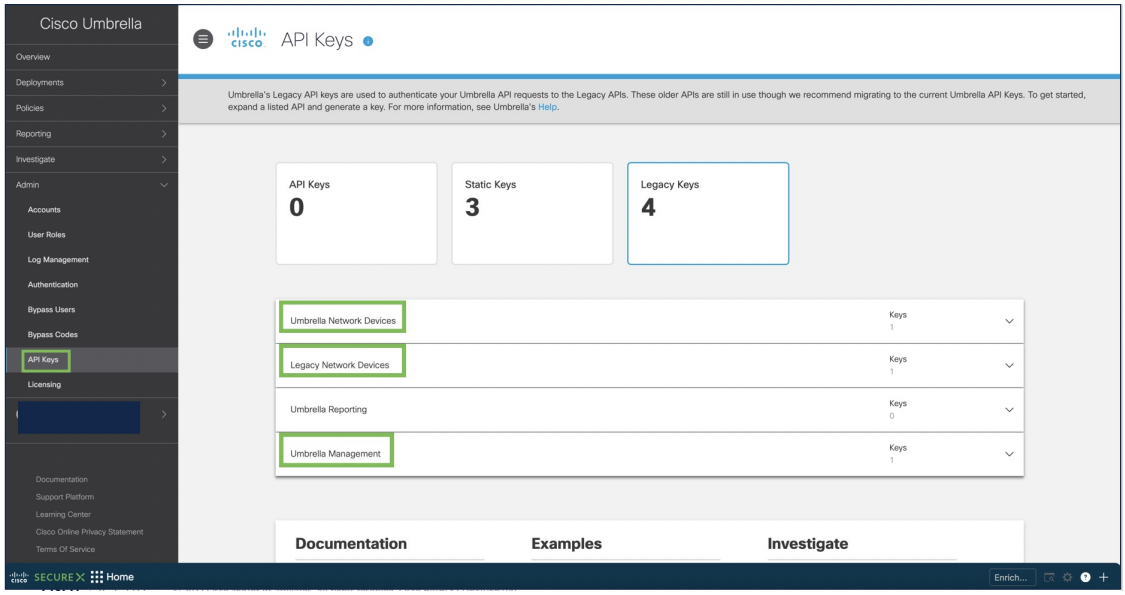

Table 79: Map Firewall Management Center Umbrella Parameters and Cisco Umbrella API Keys

Management Center Parameter	Cisco Umbrella API Keys
Network Device Key Network Device Secret	Umbrella Network Devices

Management Center Parameter	Cisco Umbrella API Keys
Legacy Network Device Token	Legacy Network Devices
Management Key Management Secret	Umbrella Management

Configure a SASE Tunnel for Umbrella

Before you begin

Ensure that you review the prerequisites and guidelines in Prerequisites for Configuring Umbrella SASE Tunnels, on page 1558 and Guidelines and Limitations for Configuring SASE Tunnels on Umbrella, on page 1556.

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site**, and click **Add**.

Step 2 Enter a name for the topology in the **Topology Name** field.

Step 3 Click the **SASE Topology** radio button and click **Create**.

Step 4 **Pre-shared Key**: This key is auto-generated according to the Umbrella PSK requirements. For a single topology, the pre-shared key is common for all threat defense spokes and Umbrella.

The device and Umbrella share this secret key, and IKEv2 uses it for authentication. If you want to configure this key, it must be between 16 and 64 characters in length, include at least one uppercase letter, one lowercase letter, one numeral, and have no special characters. Each topology must have a unique pre-shared key. If a topology has multiple tunnels, all the tunnels have the same pre-shared key.

Step 5 Choose a data center from the **Umbrella Data center** drop-down list. (Configure routing on the Firewall Threat Defense to ensure reachability of the umbrella DC from the Firewall Threat Defense.)

Step 6 Click **Add** to add a Firewall Threat Defense node.

a) Choose a Firewall Threat Defense from the **Device** drop-down list.

Only devices managed by the Firewall Management Center appear in the list. For high availability pairs, the HA pair names appear in the endpoint list.

b) Choose a static VTI interface from the **VPN Interface** drop-down list.

To create a new static VTI interface, click Add +. The **Add Virtual Tunnel Interface** dialog box appears with the following pre-populated default configurations.

- Tunnel Type is static.
- Name is <*tunnel_source interface logical name*>+ static_vti +<*tunnel ID*>. For example, outside_static_vti_2.
- Tunnel ID is auto-populated with a unique ID.
- Tunnel Source Interface is auto-populated with an interface with an 'outside' prefix.
- IPsec tunnel mode is IPv4.

- IP address is from the 169.254.x.x/30 private IP address range.

c) Enter a prefix for the local tunnel ID in the **Local Tunnel ID** field.

The prefix can have a minimum of eight characters and a maximum of 100 characters. Umbrella generates the complete tunnel ID (<*prefix*>@<*umbrella generated ID*>-umbrella.com) after the management center deploys the tunnel on Umbrella. The management center then retrieves and updates the complete tunnel ID and deploys it on the threat defense device. Each tunnel has a unique local tunnel ID.

d) Click **Save** to add the endpoint device to the topology.

You can add multiple endpoints in a SASE topology.

Step 7 Click **Next** to view the summary of the Umbrella SASE tunnel configuration.

- **Endpoints** pane: Displays the summary of the configured endpoints.
- **Encryption Settings** pane: Displays the default IKEv2 policies and IKEv2 IPsec transform sets for the topology.

Step 8 Check the **Deploy configuration on threat defense nodes** check box to trigger deployment of the network tunnels to the threat defense. This deployment occurs after the tunnels are deployed on Umbrella. Local tunnel ID is required for the threat defense deployment.

Step 9 Click **Save**.

This action:

a. Saves the topology in the management center.

b. Triggers deployment of the network tunnels to the Umbrella.

c. Triggers deployment of the network tunnels to the threat defense devices, if the option is enabled. This action commits and deploys all the updated configurations and policies, including non-VPN policies, since the last deployment on the device.

d. Opens the **Cisco Umbrella Configuration** window and displays the status of the tunnel deployment on Umbrella. For more details, see View SASE Tunnel Status, on page 1561.

What to do next

For the interesting traffic intended to flow through the SASE tunnel, configure a PBR policy with a specific match criteria to send the traffic through the VTI interface.

Ensure that you configure a PBR policy for each endpoint of the SASE topology.

View SASE Tunnel Status

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site**, and click **Add**.

Step 2 Enter a name for the topology in the **Topology Name** field.

Step 3 Click the **SASE Topology** radio button and click **Create**.

Step 4 Enter a unique **Pre-shared Key**, choose a data center, add a device, and click **Next**.

Step 5 View the summary of the Umbrella SASE Tunnel configuration and click **Save**. The **Cisco Umbrella Configuration** window appears.

You can view the topology details such as name, data center, data center IP address, start, and end time of the tunnel deployment.

You can view the deployment status of the tunnels on Umbrella. The different tunnel deployment statuses are:

- Pending: The management center hasn't pushed the configuration to Umbrella.
- Success: The management center successfully configured a tunnel on Umbrella.
- In Progress: The management center is deploying the tunnel on Umbrella.
- Failure: The management center couldn't configure a tunnel on Umbrella.

If the status appears as pending, or failure, use the transcript to troubleshoot the tunnel creation. Click the **Transcript** button to view the transcript details such as the APIs, request payload, and the response received from Umbrella.

Step 6 Click **Umbrella Dashboard** to view the network tunnels in Cisco Umbrella.

Step 7 View the Umbrella tunnel deployment status in:

- The **Notifications** page under the **Deployments** and **Tasks** tabs.
- The **Site to Site** VPN dashboard (**Overview** > **Dashboards** > **Site to Site VPN**).

  The dashboard provides a summary of the site to site VPN topologies including the Umbrella SASE topologies. You can view the tunnels between the peer devices and the status of each tunnel. You can also troubleshoot any problems with the tunnel deployment using CLI commands and packet tracer.

Monitoring the Site-to-Site VPNs

The Secure Firewall Management Center provides a snapshot of the site-to-site VPN tunnels , including SASE topology tunnels, to determine the status of the site-to-site VPN tunnels. You can view the list of tunnels between peer devices and the status of each tunnel: Active, Inactive, or No Active Data. You can filter the data in the table according to the topology, device, and status. The table in the monitoring dashboard presents live data and you can configure to refresh the data at a specified interval. The table shows the peer-to-peer, hub and spoke, and full mesh topologies for crypto map-based VPNs. The tunnel information also contains the data for the route-based VPNs or Virtual Tunnel Interfaces (VTIs).

You can use this data to:

- Identify problematic VPN tunnels and troubleshoot.
- Verify connectivity between the site to site VPN peers devices.
- Monitor the health of the VPN tunnels to provide uninterrupted VPN connectivity between sites.

For information about configuring crypto-map based site to site VPNs, see Configure a Policy-based Site-to-Site VPN, on page 1516.

For information about VTIs, see About Virtual Tunnel Interfaces, on page 1529.

For information about Firewall Threat Defense VPN monitoring and troubleshooting, see VPN Monitoring and Troubleshooting, on page 1703.

Guidelines and Limitations

- The table shows the list of site-to-site , including SASE topology, VPNs that are deployed. It does not show the tunnels that are created and not deployed.
- The table does not show the information about the backup tunnels of policy-based VPNs and backup VTIs.
- For cluster deployments, the table does not show director change in real-time data. It shows only the director information that existed when the VPN was deployed. The director change reflects in the table only after the tunnel AM redeployed after the change.

Site-to-site VPN Monitoring Dashboard

Choose **Overview** > **Dashboards** > **Site to Site VPN** to open the site to site monitoring dashboard.

The site-to-site VPN monitoring dashboard displays the following widgets for the site-to-site VPN tunnels:

- **Tunnel Status**—A table listing the tunnel status of the site-to-site VPNs , including the SASE tunnels for Umbrella, configured using the Firewall Management Center
- **Tunnel Summary**—Aggregated status of the tunnels in a donut graph.
- **Topology**—Status of tunnels summarized by topology.

Status of VPN Tunnels

The site-to-site monitoring dashboard lists the VPN tunnels in the following states:

- **Inactive**—A policy-based (crypto map-based) VPN tunnel is inactive if all the IPSec tunnels are down. A VTI or and SASE topology VPN tunnel is down if the tunnel encounters any configuration or connectivity issues.
- **Active**—In the Firewall Management Center, policy-based site-to-site VPNs are configured based on IKE policies and IPsec proposals that are assigned to VPN topologies. A policy-based VPN tunnel is in the Active state if the Firewall Management Center identifies interesting traffic through the tunnel after the deployment. An IKE tunnel is up only if a minimum of one IPsec tunnel is up.

  Route-based VPN (VTI) and SASE topology VPN tunnels do not require interesting traffic to be in the Active state. They are in the Active state if they are configured and deployed without errors.
- **No Active Data**—Policy-based and SASE topology VPN tunnels remain in the No Active Data state until there is a traffic flow event through the tunnel for the first time. The No Active Data state also lists the policy-based and route-based VPNs that have been deployed with errors.

Important Notes About Tunnel Statuses in the Firewall Management Center

- The VPN statuses in the Firewall Management Center are event-based. The Firewall Management Center does not initiate status updates. Hence, there might be mismatches between the tunnel statuses in the

dashboard and the Firewall Threat Defense. You can view the correct status in the **CLI Details** tab of the **Tunnel Status** widget.

- When a Firewall Threat Defense switches over to a secondary Firewall Threat Defense, there is a mismatch between the status of the VPN tunnels in the Firewall Management Center and the Firewall Threat Defense. When the device switches back to the primary device, the correct tunnel status appears.
- The Firewall Management Center does not update the tunnel status of Firewall Threat Defense devices earlier than 7.3 after the devices reboot. We recommend that you bring the tunnel down using the command **vpn-sessiondb logoff index** and bring it up using the packet tracer.

Tunnel Status

This table lists the site-to-site VPNs, including SASE topology VPN, configured using the Firewall Management Center. Hover over a topology and click **View** ( ) to view the following details about the topology:

- **General**—Displays more information about the nodes such as IP address and interface name.
- **CLI Details**—Displays the CLI outputs for the following commands:
  - **show crypto ipsec sa peer** <*node A/B_ip_address*>: Displays the IPsec SAs built between Node A and B.
  - **show vpn-sessiondb l2l filter ipaddress** <*node A/B_ip_address*>: Displays information about VPN sessions.

  For an extranet device, no command output appears.

  Important details about the IKE and IPsec sessions, derived from the above command outputs, appear in a summarized, and user-friendly format. You can view the details of both nodes at a time. The icon next to the node name specifies the authentication type: preshared key or client certificate. The details include IKE statistics per tunnel and IPsec SA statistics as shown below:

Figure 428: Tunnel Status > View > CLI Details

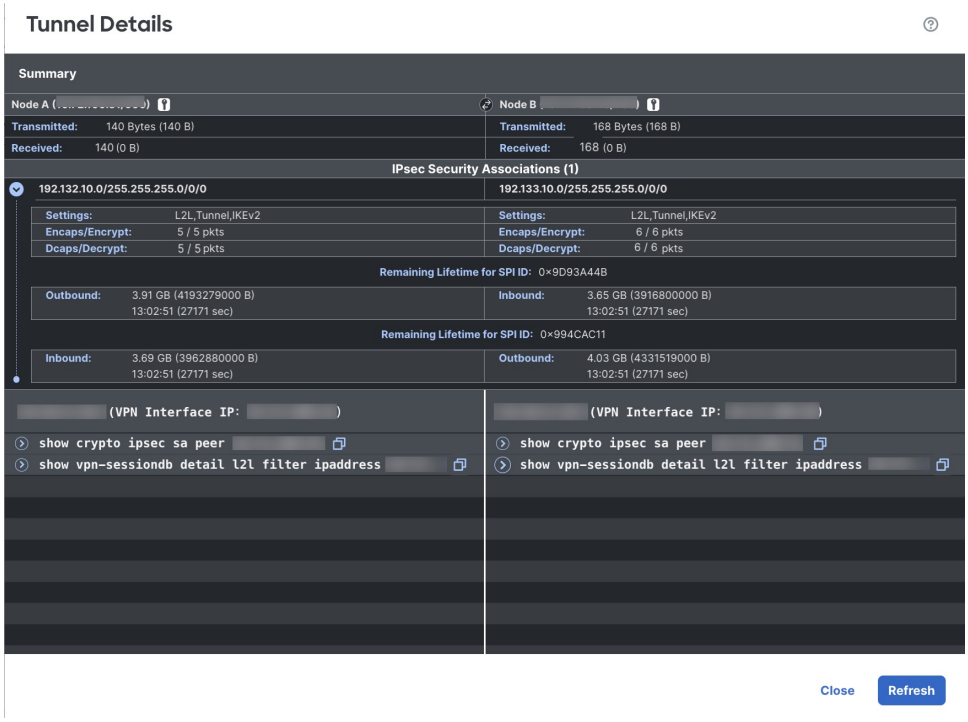

- **Packet Tracer**—Use packet tracer to troubleshoot the threat defense VPN tunnels.

Packet Tracer

Packet tracer allows you to troubleshoot VPN tunnels between two threat defense devices. You can check if the VPN connection between device A and device B is up. This tool injects a packet into the device and tracks the packet flow from the ingress to the egress ports. The tool simulates traffic after you configure the ingress interfaces of the devices along with the protected networks. Packet tracer evaluates the packet against modules such as flow and route lookups, ACLs, protocol inspection, NAT, and QoS.

Figure 429: Packet Tracer

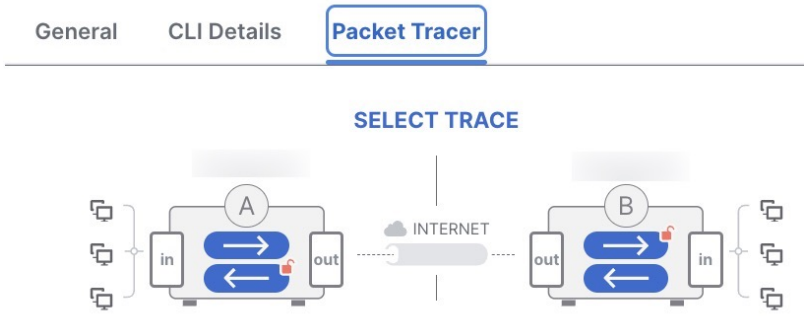

For each device, the tool runs an encrypted trace and a decrypted trace (packet is treated as decrypted VPN traffic). You can run four different traces between the ingress and egress ports of the devices. Click the individual encrypt and decrypt options to enable or disable the trace.

When you run the trace, the tool executes the trace sequentially in the following order:

1. Encrypted trace of A.
2. Decrypted trace of B.
3. Encrypted trace of B.
4. Decrypted trace of A.

After the trace completes, you can view the output of the trace with the results of each module.

Note You cannot run a decrypt trace for route-based (VTI-based) VPNs.

To run the Packet Tracer:

1. Click **See Detailed Config** to view the VPN interface name, VPN interface IP address, VTI interface name, and the VTI Interface IP address.
2. (Optional) Choose a protocol from the **Protocol** drop-down list. You can choose ICMP/8/0, TCP, or UDP.

   ICMP/8/0 is the default option. If you choose ICMP/8/0, 8 indicates the ICMP type as Echo Request and 0 indicates the ICMP code. If you choose TCP or UDP, choose the destination port from the **Destination Port** drop-down list. The range is from 0 to 65535.
3. Choose the ingress interface for both the devices on which to trace the packet from the **Ingress Interface** drop-down lists.
4. Enter an IP address from the same subnet as the ingress interface in the **Protected Network IP Address** fields.
5. Click **Trace Now**.

   After you initiate the trace, you can view if the trace is successful or not for each module. If the tunnel is down, the path appears in red. If the tunnel is up, the path appears in green. If a tunnel is down, click **Re-trace** to run the tool again. For a crypto-map based VPN, when the tunnel is inactive with no interesting traffic, the initial trace can be red. Click **Re-trace** to run the trace again.

Figure 430: Packet Tracer after a Successful Trace

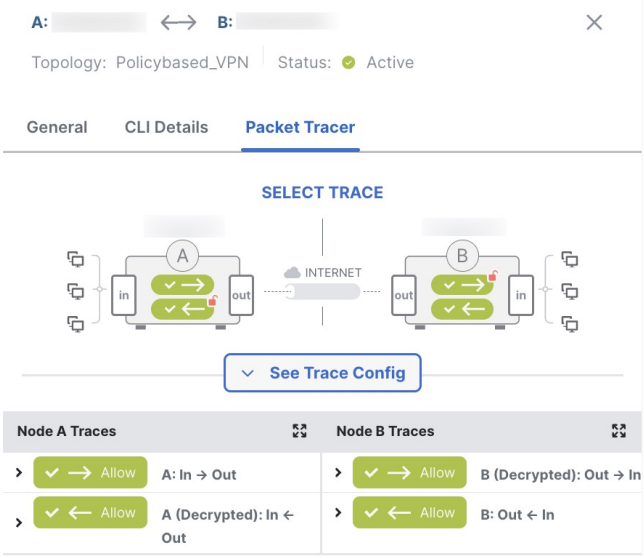

Extranet Nodes: You can initiate a packet trace for VPN tunnels with one node as an extranet. For an extranet node, you cannot choose the ingress interface. The remaining steps for the packet trace are the same. You can't run trace on the extranet side.

For example, if Node A is a managed threat defense and Node B is an extranet:

- Configure the ingress interface for node A.
- Configure the protected network for Node A and B.
- Click **Trace Now**. The traces appear for Node A and not for Node B.

Automatic Data Refresh

The site to site VPN data in the table refreshes periodically. You can configure the refresh interval of the VPN monitoring data at a specific interval or turn the automatic data refresh off.

Click the **Refresh** interval drop-down to select from the available intervals to refresh the data in the table.

Figure 431: Refresh the Tunnel Data

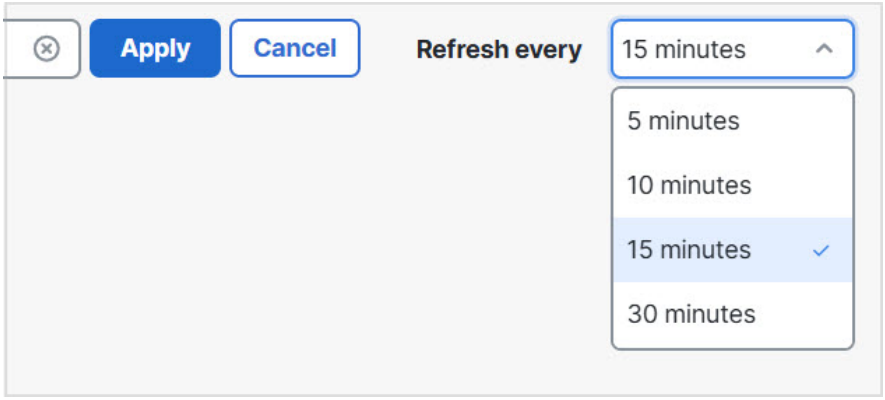

Click **Pause** to stop the automatic data refresh for as long as you want. You can click the same button to resume refreshing the tunnel data.

Figure 432: Pause the Periodic Data Refresh

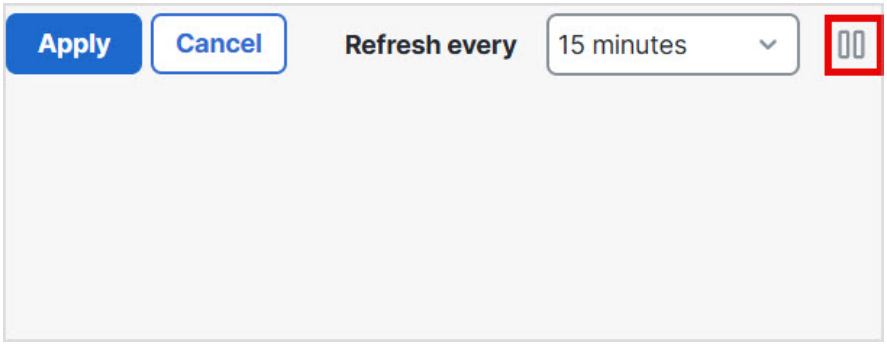

Filter and Sort the Site to Site VPN Monitoring Data

You can filter and view the data in the VPN monitoring table by topology, device, and status.

For example, you can view the tunnels that are in the Down state in a specific topology.

Click within the filter box to choose the filter criteria and then specify the values to filter.

Figure 433: Filter the Tunnel Data

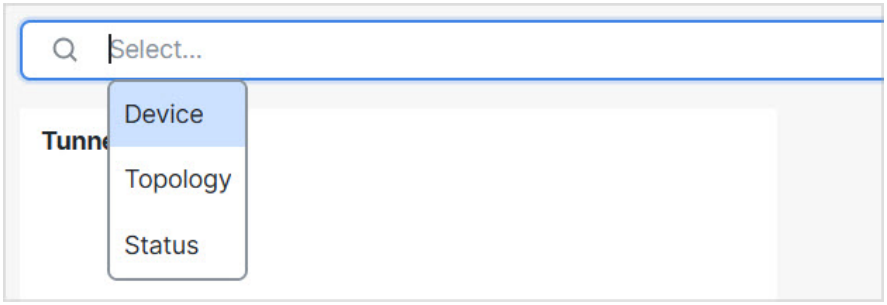

You can use multiple filtering criteria to view the data based on your requirement.

For example, you can choose to view only the tunnels that are in the Up and Down states, and ignore the ones in the Unknown state.

Figure 434: Example: Filter Tunnel Data

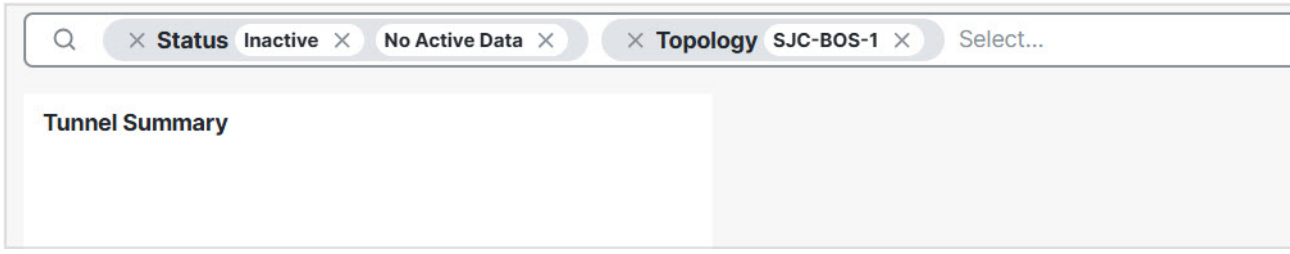

Sort the data—To sort the data by a column, click the column heading.

Related Topics

About Site-to-Site VPN, on page 1511
About Virtual Tunnel Interfaces, on page 1529

History for Site-to-Site VPN

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
View Virtual Tunnel Information	7.4.1	Any	You can view details of dynamic and static VTIs of route-based VPNs on the device. For every VPN topology, you can also view details of all the virtual access interfaces of the dynamic VTIs. New/Modified screens: **Device** > **Device Management** > **Edit a device** > **Interfaces** and click the **Virtual Tunnels** tab.
IPsec flow offload	7.4	Any	IPsec flow offload is automatically enabled on the VTI loopback interface on Secure Firewall 3100 and Secure Firewall 4200 devices.
Umbrella SASE Topology	7.3	Any	You can configure an Umbrella SASE topology and deploy IPsec IKEv2 tunnels between a threat defense device and Umbrella. This tunnel forwards all internet-bound traffic to the Umbrella Secure Internet Gateway (SIG) for inspection and filtering.
Support for Dynamic Virtual Tunnel Interface	7.3	Any	You can create a dynamic VTI and use it to configure a route-based site-to-site VPN in a hub and spoke topology.
EIGRP IPv4 support for VTI	7.3	Any	Static and dynamic VTI interfaces support EIGRP IPv4 routing protocol.
OSPFv2/v3 IPv4/v6 support for VTI	7.3	Any	Static and dynamic VTI interfaces support OSPFv2/v3 IPv4/v6 routing protocol.
Packet Tracer in Site to Site VPN Monitoring Dashboard	7.3	Any	Use the packet tracer tool in the site-to-site VPN monitoring dashboard to troubleshoot the threat defense VPN tunnels. New/Modified screens: **Overview** > **Dashboards** > **Site to Site VPN**
Remote Access VPN Dashboard	7.3	Any	Use the Remote Access VPN dashboard to monitor real-time data from active remote access VPN sessions on the devices. New/Modified screens: **Overview** > **Dashboards** > **Site to Site VPNRemote Access VPN**

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
IPsec flow offload	7.2	Any	On the Secure Firewall 3100, IPsec flows are offloaded by default. After the initial setup of an IPsec site-to-site VPN or remote access VPN security association (SA), IPsec connections are offloaded to the field-programmable gate array (FPGA) in the device, which should improve device performance. You can change the configuration using FlexConfig and the **flow-offload-ipsec** command.
Site-to-Site VPN Filter	7.1	Any	You can control site-to-site VPN traffic by using an access control policy.
Local tunnel ID support	7.1	Any	For each endpoint on a site-to-site VPN, you can configure a unique tunnel ID to be shared with the peers.
Multiple IKE Policy Support	7.1	Any	You can add multiple IKEv1 and IKEv2 policy objects for each endpoint.
Site-to-Site VPN Monitoring Dashboard	7.1	Any	Use the Site-to-Site VPN Monitoring dashboard to view and monitor the status of site-to-site VPN tunnels.
Backup virtual tunnel interfaces (VTI) for route-based site-to-site VPN.	7.0	Any	When you configure a site-to-site VPN that uses virtual tunnel interfaces, you can select a backup VTI for the tunnel. Specifying a backup VTI provides resiliency, so that if the primary connection goes down, the backup connection might still be functional. For example, you could point the primary VTI to the endpoint of one service provider, and the backup VTI to the endpoint of a different service provider. You can add a backup VTI in the site-to-site VPN wizard by selecting route-based as the VPN type for a point-to-point connection.
Enhance the number of VTI from 100 per interface to 1024 per device	7.0	Any	Support for maximum number of VTIs is enhanced from 100 per physical interface to 1024 VTIs per device.
IPv6 Support	7.0	Any	You can configure IPv6 addressed VTIs. While only static IPv6 address is supported as the tunnel source and destination, IPv6 BGP isn't supported over VTI.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Removal and deprecation of weak ciphers	6.7	Any	Support has been removed for less secure ciphers. We recommend that you update your VPN configuration before you upgrade to Firewall Threat Defense 6.70 to supported DH and encryption algorithms to ensure the VPN works correctly. Update your IKE proposals and IPSec policies to match the ones supported in Firewall Threat Defense 6.70 and then deploy the configuration changes. The following less secure ciphers have been removed or deprecated in Firewall Threat Defense 6.70 onwards: • **Diffie-Hellman GROUP 5** is deprecated for IKEv1 and removed for IKEv2 • Diffie-Hellman groups 2 and 24 have been removed. • **Encryption algorithms**: 3DES, AES-GMAC, AES-GMAC-192, AES-GMAC-256 have been removed. **Note** **DES** is supported in evaluation mode or for users who do not satisfy export controls for strong encryption. **NULL** is removed in IKEv2 policy, but supported in both IKEv1 and IKEv2 IPsec transform-sets.
Dynamic RRI support	6.7	Any	Dynamic Reverse Route Injection is supported with IKEv2 based static crypto maps.
Backup peer for site-to-site VPN	6.6	Any	You can use the Firewall Management Center to add a backup peer to a site-to-site VPN connection. For example, if you have two ISPs, you can configure the VPN connection to fail over to the backup ISP if the connection to the first ISP becomes unavailable. New/modified pages: **Devices** > **VPN** > **Site to Site**. When adding or editing a point to point or hub and spoke FTD VPN topology to add an endpoint, the **IP Address** field supports comma-separated backup peers.

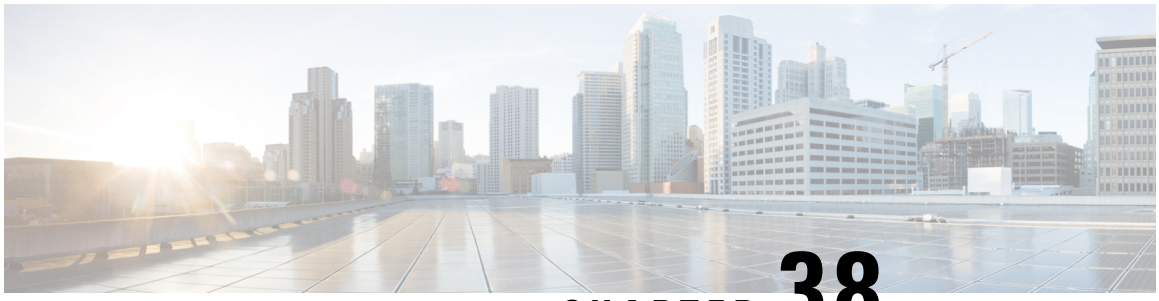

CHAPTER 38

Remote Access VPN

Remote Access virtual private network (VPN) allows individual users to connect to your network from a remote location using a computer or other supported devices connected to the Internet. This allows mobile workers to connect from their home networks or a public Wi-Fi network, for example.

The following topics explain how to configure remote access VPN for your network.

- Remote Access VPN Overview, on page 1573
- License Requirements for Remote Access VPN, on page 1579
- Requirements and Prerequisites for Remote Access VPN, on page 1579
- Guidelines and Limitations for Remote Access VPNs, on page 1580
- Configuring a New Remote Access VPN Connection, on page 1583
- Create a Copy of an Existing Remote Access VPN Policy, on page 1593
- Set Target Devices for a Remote Access VPN Policy, on page 1593
- Associate Local Realm with Remote Access VPN Policy, on page 1594
- Additional Remote Access VPN Configurations, on page 1594
- Customizing Remote Access VPN AAA Settings, on page 1650
- Advanced Secure Client Configurations, on page 1670
- Remote Access VPN Examples, on page 1679
- History for Remote Access VPNs, on page 1684

Remote Access VPN Overview

Secure Firewall Threat Defense provides secure gateway capabilities that support remote access SSL and IPsec-IKEv2 VPNs. The full tunnel client, Secure Client, provides secure SSL and IPsec-IKEv2 connections to the security gateway for remote users. When the client negotiates an SSL VPN connection with Firewall Threat Defense, it connects using Transport Layer Security (TLS) or Datagram Transport Layer Security (DTLS)

Secure Client is the only client supported on endpoint devices for remote VPN connectivity to Firewall Threat Defense devices. The client gives remote users the benefits of an SSL or IPsec-IKEv2 VPN client without the need for network administrators to install and configure clients on remote computers. The Secure Client for Windows, Mac, and Linux is deployed from the secure gateway upon connectivity. The Secure Client apps for Apple iOS and Android devices are installed from the platform app store.

Use the Remote Access VPN policy wizard to set up SSL and IPsec-IKEv2 remote access VPNs with basic capabilities. Then, enhance the policy configuration as you want and deploy it to your Firewall Threat Defense secure gateway devices.

Remote Access VPN Features

The following table describes the features of Secure Firewall Threat Defense remote access VPN:

Table 80: Remote access VPN features

	Description
Secure Firewall Threat Defense remote access VPN features	• SSL and IPsec-IKEv2 remote access using the Secure Client. • Secure Firewall Management Center supports all combinations such as IPv6 over an IPv4 tunnel. • Configuration support on both Firewall Management Center and Firewall Device Manager. Device-specific overrides. • Support for both Secure Firewall Management Center and Firewall Threat Defense HA environments. • Support for multiple interfaces and multiple AAA servers. • Rapid Threat Containment support using RADIUS CoA or RADIUS dynamic authorization. • Support for DTLS v1.2 protocol with Cisco Secure Client version 4.7 or higher. • Secure Client modules support for additional security services for remote access VPN connections. • VPN load balancing.

	Description
AAA features	• Server authentication using self-signed or CA-signed identity certificates. • AAA username and password-based remote authentication using RADIUS server or LDAP or AD. • RADIUS group and user authorization attributes, and RADIUS accounting. • Double authentication support using an additional AAA server for secondary authentication. • NGFW Access Control integration using VPN Identity. • LDAP or AD authorization attributes using Secure Firewall Management Center web interface. • Support for single sign-on using SAML 2.0. • Support for multiple identity provider trustpoints with Microsoft Azure that can have multiple applications for the same Entity ID, but a unique identity certificate. • Restrict remote access VPN connections based on their geolocations.
VPN tunneling features	• Address assignment. • Split tunneling. • Split DNS. • Client Firewall ACLs. • Session Timeouts for maximum connect and idle time.
Remote access VPN monitoring features	• New VPN Dashboard Widget showing VPN users by various characteristics such as duration and client application. • Remote access VPN events including authentication information such as username and OS platform. • Tunnel statistics available using the Firewall Threat Defense Unified CLI.

Secure Client Components

Secure Client Deployment

Your remote access VPN policy can include the Secure Client Image and the Secure Client Profile for distribution to connecting endpoints. Or, the client software can be distributed using other methods. See the *Deploy Cisco Secure Client* chapter in the Cisco Secure Client (including AnyConnect) Administrator Guide, Release 5.

Without a previously installed client, remote users enter the IP address in their browser of an interface configured to accept SSL or IPsec-IKEv2 VPN connections. Unless the security appliance is configured to redirect http:// requests to https://, remote users must enter the URL in the form https://*address*. After the user enters the URL, the browser connects to that interface and displays the login screen.

After a user logs in, if the secure gateway identifies the user as requiring the VPN client, it downloads the client that matches the operating system of the remote computer. After downloading, the client installs and configures itself, establishes a secure connection, and either remains or uninstalls itself (depending on the security appliance configuration) when the connection stops. In the case of a previously installed client, after login, the Firewall Threat Defense security gateway examines the client version and upgrades it as necessary.

Secure Client Operation

When the client negotiates a connection with the security appliance, the client connects using Transport Layer Security (TLS), and optionally, Datagram Transport Layer Security (DTLS). DTLS avoids latency and bandwidth problems associated with some SSL connections and improves the performance of real-time applications that are sensitive to packet delays.

When an IPsec-IKEv2 VPN client initiates a connection to the secure gateway, negotiation consists of authenticating the device through Internet Key Exchange (IKE), followed by user authentication using IKE Extended Authentication (Xauth). The group profile is pushed to the VPN client and an IPsec security association (SA) is created to complete the VPN.

Secure Client Profile and Editor

The Secure Client Profile is a group of configuration parameters, stored in an XML file that the VPN client uses to configure its operation and appearance. These parameters (XML tags) include the names and addresses of host computers and settings to enable more client features.

You can configure a profile using the Secure Client Profile Editor. This editor is a convenient GUI-based configuration tool that is available as part of the Secure Client software package. It is an independent program that you run outside of the Firewall Management Center.

Remote Access VPN Authentication

Remote Access VPN Server Authentication

Secure Firewall Threat Defense secure gateways always use certificates to identify and authenticate themselves to the VPN client endpoint.

While you use the Remote Access VPN Policy Wizard, you can enroll the selected certificate on the targeted Firewall Threat Defense device. In the wizard, under **Access & Certificate** phase, select “Enroll the selected certificate object on the target devices” option. The certificate enrollment gets automatically initiated on the specified devices. As you complete the remote access VPN policy configuration, you can view the status of

the enrolled certificate under the device certificate homepage. The status provides a clear standing as to whether the certificate enrollment was successful or not. Your remote access VPN policy configuration is now fully completed and ready for deployment.

Obtaining a certificate for the secure gateway, also known as PKI enrollment, is explained in Certificates, on page 1485. This chapter contains a full description of configuring, enrolling, and maintaining gateway certificates.

Remote Access VPN Client AAA

For both SSL and IPsec-IKEv2, remote user authentication is done using usernames and passwords only, certificates only, or both.

Note If you are using client certificates in your deployment, they must be added to your client's platform independent of the Secure Firewall Threat Defense or Secure Firewall Management Center. Facilities such as SCEP or CA Services are not provided to populate your clients with certificates.

AAA servers enable managed devices acting as secure gateways to determine who a user is (authentication), what the user is permitted to do (authorization), and what the user did (accounting). Some examples of the AAA servers are RADIUS, LDAP/AD, TACACS+, and Kerberos. For Remote Access VPN on Firewall Threat Defense devices, AD, LDAP, and RADIUS AAA servers are supported for authentication.

Refer to the section Understanding Policy Enforcement of Permissions and Attributes to understand more about remote access VPN authorization.

Before you add or edit the remote access VPN policy, you must configure the Realm and RADIUS server groups you want to specify. For more information, see Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426 and Add a RADIUS Server Group, on page 1364.

Without DNS configured, the device cannot resolve AAA server names, named URLs, and CA Servers with FQDN or Hostnames, it can only resolve IP addresses.

The login information provided by a remote user is validated by an LDAP or AD realm or a RADIUS server group. These entities are integrated with the Secure Firewall Threat Defense secure gateway.

Note If users authenticate with remote access VPN using Active Directory as the authentication source, users must log in using their username; the format `domain\username` or `username@domain` fails. (Active Directory refers to this username as the *logon name* or sometimes as `sAMAccountName`.) For more information, see User Naming Attributes on MSDN.

If you use RADIUS to authenticate, users can log in with any of the preceding formats.

Once authenticated via a VPN connection, the remote user takes on a *VPN Identity*. This VPN Identity is used by *identity policies* on the Secure Firewall Threat Defense secure gateway to recognize and filter network traffic belonging to that remote user.

Identity policies are associated with access control policies, which determine who has access to network resources. It is in this way that the remote user blocked or allowed to access your network resources.

For more information, see the About Identity Policies, on page 2553 and Access Control Policies, on page 1745 sections.

Related Topics

Configure AAA Settings for Remote Access VPN, on page 1597

Understanding Policy Enforcement of Permissions and Attributes

The Secure Firewall Threat Defense device supports applying user authorization attributes (also called user entitlements or permissions) to VPN connections from an external authentication server and/or authorization AAA server (RADIUS) or from a group policy on the Firewall Threat Defense device. If the Firewall Threat Defense device receives attributes from the external AAA server that conflicts with those configured on the group policy, then attributes from the AAA server always take the precedence.

The Firewall Threat Defense device applies attributes in the following order:

1. **User attributes on the external AAA server**—The server returns these attributes after successful user authentication and/or authorization.
2. **Group policy configured on the Firepower Threat Defense device**—If a RADIUS server returns the value of the RADIUS Class attribute IETF-Class-25 (OU= group-policy) for the user, the Firewall Threat Defense device places the user in the group policy of the same name and enforces any attributes in the group policy that are not returned by the server.
3. **Group policy assigned by the Connection Profile (also known as Tunnel Group)**—The Connection Profile has the preliminary settings for the connection, and includes a default group policy applied to the user before authentication.

Note The Firewall Threat Defense device does not support inheriting system default attributes from the default group policy, *DfltGrpPolicy*. The attributes on the group policy assigned to the connection profile are used for the user session, if they are not overridden by user attributes or the group policy from the AAA server as indicated above.

Related Topics

Configure AAA Settings for Remote Access VPN, on page 1597

Understanding AAA Server Connectivity

LDAP, AD, and RADIUS AAA servers must be reachable from the Firewall Threat Defense device for your intended purposes: user-identity handling only, VPN authentication only, or both activities. AAA servers are used in remote access VPN for the following activities:

- **User-identity handling**— the servers must be reachable over the Management interface.

  On the Firewall Threat Defense, the Management interface has a separate routing process and configuration from data interfaces.

- **VPN authentication**—the servers must be reachable over a data interface or the Management interface.

  To use the Management interface, you must explicitly select Management as the source interface. Other management-only interfaces cannot be used.

To use the same AAA servers for both activities, we recommend specifying the Management interface as the source interface.

For more information about various interfaces, see Regular Firewall Interfaces, on page 817.

After deployment, use the following CLI commands to monitor and troubleshoot AAA server connectivity from the Firewall Threat Defense device:

- **show aaa-server** to display AAA server statistics.
- **show network** and **show network-static-routes** to view the Management interface default route and static routes.
- **show route** to view data traffic routing table entries.
- **ping system** and **traceroute system** to verify the path to the AAA server through the Management interface.
- **ping interface** *ifname* and **traceroute** *destination* to verify the path to the AAA server through the data interfaces.
- **test aaa-server authentication** and **test aaa-server authorization** to test authentication and authorization on the AAA server.
- **clear aaa-server statistics** *groupname* or **clear aaa-server statistics protocol** *protocol* to clear AAA server statistics by group or protocol.
- **aaa-server** *groupname* **active host** *hostname* to activate a failed AAA server, or **aaa-server** *groupname* **fail host** *hostname* to fail a AAA server.
- **debug ldap level**, **debug aaa authentication**, **debug aaa authorization**, and **debug aaa accounting**.

License Requirements for Remote Access VPN

Threat Defense License

Firewall Threat Defense remote access VPN requires Strong Encryption and one of the following licenses for Secure Client:

- Secure Client Advantage
- Secure Client Premier
- Secure Client VPN Only

Requirements and Prerequisites for Remote Access VPN

Model Support

Threat Defense

Supported Domains

Any

User Roles

Admin

Guidelines and Limitations for Remote Access VPNs

Remote Access VPN Policy Configuration

- You can add a new remote access VPN policy only by using the wizard. You must proceed through the entire wizard to create a new policy; the policy will not be saved if you cancel before completing the wizard.
- Two users must **not** edit a remote access VPN policy at the same time; however, the web interface does not prevent simultaneous editing. If this occurs, the last saved configuration persists.
- Moving a Secure Firewall Threat Defense device from one domain to another domain is not possible if remote access VPN policy is assigned to that device.
- Remote access VPN does not support SSL while using ECMP. We recommend that you use IPsec-IKEv2.
- Firepower 9300 and 4100 series in cluster mode do not support remote access VPN configuration.
- Remote access VPN connectivity could fail if there is a misconfigured Firewall Threat Defense NAT rule.
- If you are using DHCP to provide IP addresses to the client, and the client cannot obtain an address, check the NAT rules. Any NAT rule that applies to the RA VPN network should include the route lookup option. Route lookup can help ensure the DHCP requests are sent to the DHCP server through an appropriate interface.
- Whenever IKE ports 500/4500 or SSL port 443 is in use or when there are some PAT translations that are active, the Secure Client IPSec-IKEv2 or SSL remote access VPN cannot be configured on the same port as it fails to start the service on those ports. These ports must not be used on the Firewall Threat Defense device before configuring remote access VPN policy.
- While configuring remote access VPNs using the wizard, you can create in-line certificate enrollment objects, but you cannot use them to install the identity certificate. Certificate enrollment objects are used for generating the identity certificate on the Firewall Threat Defense device being configured as the remote access VPN gateway. Install the identity certificate on the device before deploying the remote access VPN policy to the device.

  For more information about how to install the identity certificate based on the certificate enrollment object, see The Object Manager, on page 1354.
- The ECMP zone interfaces can be used in remote access VPN with IPsec enabled.
- The ECMP zone interfaces cannot be used in remote access VPN with SSL enabled. Deployment of remote access VPN (SSL enabled) configuration fails if all the remote access VPN interfaces that belong to security zones or interface groups also belong to one or more ECMP zones. However, if only some of the remote access VPN interfaces belonging to the security zones or interface groups also belongs to one or more ECMP zones, deployment of the remote access VPN configuration succeeds excluding those interfaces.
- After you change the remote access VPN policy configurations, re-deploy the changes to the Firewall Threat Defense devices. The time it takes to deploy configuration changes depends on multiple factors

such as complexity of the policies and rules, type and volume of configurations you send to the device, and memory and device model. Before deploying remote access VPN policy changes, review the Best Practices for Deploying Configuration Changes, on page 258.

- Issuing commands such as **curl** against the RA VPN headend is not directly supported, and might not have desirable results. For example, the headend does not respond to HTTP HEAD requests.
- Firewall Threat Defense does not accept remote access VPN sessions if the third-party client sends a null user agent.
- You can configure browser proxy using FlexConfig.
- In a remote branch deployment (RBD) with High Availability, when you break a High Availability pair, the RA VPN policy assignment and VPN configurations will be removed on the RBD WAN interface of the standby unit.
- If a realm has more than the 16 directories, only the first 16 directories in the realm are deployed during an RA VPN connection.
- When **Bypass Access Control policy for decrypted traffic** is disabled on a Firewall Threat Defense and ISE issues a DACL, traffic is first matched against the ACP: if permitted, it proceeds to DACL checking; if denied, it bypasses DACL matching.

Concurrent VPN Sessions Capacity Planning (Firewall Threat Defense Virtual Models)

The maximum concurrent VPN sessions are governed by the installed Firewall Threat Defense Virtual smart-licensed entitlement tier, and enforced via a rate limiter. There is a maximum limit to the number of concurrent remote access VPN sessions allowed on a device based on the licensed device model. This limit is designed so that system performance does not degrade to unacceptable levels. Use these limits for capacity planning.

Device Model	Maximum Concurrent Remote Access VPN Sessions
Firewall Threat Defense Virtual5	50
Firewall Threat Defense Virtual10	250
Firewall Threat Defense Virtual20	250
Firewall Threat Defense Virtual30	250
Firewall Threat Defense Virtual50	750
Firewall Threat Defense Virtual100	10,000

Concurrent VPN Sessions Capacity Planning (Hardware Models)

The maximum concurrent VPN sessions are governed by platform-specific limits and have no dependency on the license. There is a maximum limit to the number of concurrent remote access VPN sessions allowed on a device based on the device model. This limit is designed so that system performance does not degrade to unacceptable levels. Use these limits for capacity planning.

Device Model	Maximum Concurrent Remote Access VPN Sessions
Firepower 1010	75

Device Model	Maximum Concurrent Remote Access VPN Sessions
Firepower 1120	150
Firepower 1140	400
Firepower 2110	1500
Firepower 2120	3500
Firepower 2130	7500
Firepower 2140	10,000
Secure Firewall 3110	3000
Secure Firewall 3120	6000
Secure Firewall 3130	15,000
Secure Firewall 3140	20,000
Firepower 4100, all models	10,000
Firepower 9300 appliance, all models	20,000
ISA 3000	25

For capacity of other hardware models, contact your sales representative.

Note The Firewall Threat Defense device denies the VPN connections once the maximum session limit per platform is reached. The connection is denied with a syslog message. Refer the syslog messages %ASA-4-113029 and %ASA-4-113038 in the syslog messaging guide. For more information, see Cisco Secure Firewall ASA Series Syslog Messages.

Controlling Cipher Usage for VPN

To prevent use of ciphers greater than DES, pre-deployment checks are available at the following locations in the Firewall Management Center:

Devices > **Platform Settings**, click **Edit**and then **SSL**.

Devices > **VPN** > **Remote Access** > **Edit** > **Advanced** > **IPsec**.

For more information about SSL settings and IPsec, see SSL , on page 980 and Configure Remote Access VPN IPsec/IKEv2 Parameters, on page 1628.

Authentication, Authorization, and Accounting

Configure DNS on each device in the topology in to use remote access VPN. Without DNS, the device cannot resolve AAA server names, named URLs, and CA Servers with FQDN or Hostnames; it can only resolve IP addresses.

You can configure DNS using the **Platform Settings**. For more information, see DNS, on page 949 and DNS Server Group, on page 1384.

Client Certificates

If you are using client certificates in your deployment, they must be added to your client's platform independent of the Secure Firewall Threat Defense or Secure Firewall Management Center. Facilities such as SCEP or CA Services are not provided to populate your clients with certificates.

Unsupported Features of Secure Client

The only supported VPN client is the Cisco Secure Client. No other clients or native VPNs are supported. Clientless VPN is not supported for VPN connectivity; it is only used to deploy the Secure Client using a web browser.

Note Using multiple Secure Client packages on Firewall Threat Defense devices can increase memory usage and affect the device's performance. We recommend that you do not use multiple Secure Client packages on low-end threat defense devices to avoid continuous reboots because of memory exhaustion.

The following Secure Client features are not supported when connecting to a Firewall Threat Defense secure gateway:

- TACACS, Kerberos (KCD Authentication and RSA SDI), and SDI.

Configuring a New Remote Access VPN Connection

This section provides instructions to configure a new remote access VPN policy with Secure Firewall Threat Defense devices as VPN gateways and Cisco Secure Client as the VPN client.

Step	Do This	More Info
1	Review the guidelines and prerequisites.	Guidelines and Limitations for Remote Access VPNs, on page 1580 Prerequisites for Configuring Remote Access VPN, on page 1584
2	Create a new remote access VPN policy using the wizard.	Create a New Remote Access VPN Policy, on page 1584
3	Update the access control policy deployed on the device.	Update the Access Control Policy on the Secure Firewall Threat Defense Device, on page 1587
4	(Optional) Configure a NAT exemption rule if NAT is configured on the device.	(Optional) Configure NAT Exemption, on page 1588
5	Configure DNS.	Configure DNS, on page 1589
6	Add Secure Client Profile.	Add Secure Client Profile XML File, on page 1589
7	Deploy the remote access VPN policy.	Deploy Configuration Changes, on page 259

Step	Do This	More Info
8	(Optional) Verify the remote access VPN policy configuration.	Verify the Configuration, on page 1592

Prerequisites for Configuring Remote Access VPN

- Deploy Secure Firewall Threat Defense devices and configure Secure Firewall Management Center to manage the device with required licenses with export-controlled features enabled. For more information, see VPN Licensing, on page 1501.
- Configure the certificate enrollment object that is used to obtain the identity certificate for each Firewall Threat Defense device that act as a remote access VPN gateway.
- Configure any AD or LDAP realms being used by remote access VPN policies.
- During migration of FTD with remote access VPN, the realm (LDAP, AD or even local) used in remote access VPN should be preconfigured on cdFMC before you migrate the remote access VPN.
- Ensure that the AAA Server is reachable from the Firewall Threat Defense device for the remote access VPN configuration to work. Configure routing (at **Devices** > **Device Management**, click **Edit Device** and then **Routing**) to ensure connectivity to the AAA servers.

  For remote access VPN double authentication, ensure that both the primary and secondary authentication servers are reachable from the Firewall Threat Defense device for the double authentication configuration to work.
- Purchase and enable one of the following Cisco Secure Client licenses: Secure Client Advantage, Secure Client Premier, or Secure Client VPN Only to enable the Firewall Threat Defense remote access VPN.
- Download the latest Secure Client image files from Cisco Software Download Center.

  On your Secure Firewall Management Center web interface, go to **Objects** > **Object Management** > **VPN** > **Secure Client File** and add the new Secure Client image files.
- Create a security zone or interface group that contains the network interfaces that users will access for VPN connections. See Interface, on page 1398.
- Download the Secure Client Profile Editor from Cisco Software Download Center to create the Secure Client profile. You can use the standalone profile editor to create a new or modify an existing Secure Client profile.

Create a New Remote Access VPN Policy

The Remote Access VPN Policy Wizard guides you to quickly and easily set up remote access VPNs with basic capabilities. You can further enhance the policy configuration by specifying additional attributes as you want and deploy it to your Secure Firewall Threat Defense secure gateway devices.

Before you begin

- Ensure that you complete all the prerequisites listed in Prerequisites for Configuring Remote Access VPN, on page 1584.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Add** to create a new remote access VPN policy with basic policy configuration, using the Remote Access VPN Policy wizard.

You must proceed through the entire wizard to create a new policy; the policy is not saved if you cancel before you complete the wizard.

Step 3 Select the target devices and protocols.

The Firewall Threat Defense devices that you select here functions as your remote access VPN gateways for the VPN client users.

You can select Firewall Threat Defense devices when you create a remote access VPN policy or change them later. See Set Target Devices for a Remote Access VPN Policy, on page 1593.

You can select **SSL** or **IPSec-IKEv2**, or both the VPN protocols. Firewall Threat Defense supports both the protocols to establish secure connections over a public network through VPN tunnels.

Note
Firewall Threat Defense does not support IPSec tunnels with NULL encryption. If you have selected IPSec-IKEv2, make sure that you do not choose NULL encryption for IPSec IKEv2 proposal. See Configure IKEv2 IPsec Proposal Objects, on page 1474.

For SSL settings, see SSL , on page 980.

Step 4 Click **Next**.

Step 5 Configure the **Connection Profile** and **Group Policy** settings.

A connection profile specifies a set of parameters that define how the remote users connect to the VPN device. The parameters include settings and attributes for authentication, address assignments to VPN clients, and group policies. Firewall Threat Defense device provides a default connection profile named *DefaultWEBVPNGroup* when you configure a remote access VPN policy.

For more information, see Configure Connection Profile Settings, on page 1594.

Step 6 Configure the **Authentication, Authorization & Accounting** settings.

For information about configuring,

- AAA settings, see Configure AAA Settings for Remote Access VPN, on page 1597
- LDAP attribute maps, see Configuring LDAP Attribute Mapping, on page 1619
- SAML 2.0 single sign-on authentication, see Configuring a SAML Single Sign-On Authentication, on page 1668

Step 7 Configure the **Client Address Assignment** settings.

Client IP address can be assigned from AAA server, DHCP server and IP address pools. When multiple options are selected, IP address assignment is done in the order of AAA server, DHCP server, and IP address pool. Assignment of client IP addresses from the AAA server is supported only for realm and RADIUS authorization. Ensure that realm or RADIUS server is configured to provide client IP address.

Step 8 Configure the **Group Policy** settings.

A group policy is a set of attribute and value pairs, stored in a group policy object, that define the remote access VPN experience for VPN users. You configure attributes such as user authorization profile, IP addresses, Secure Client settings, VLAN mapping, and user session settings and so on using the group policy. The RADIUS authorization server assigns the group policy, or it is obtained from the current connection profile. For more information, see Configuring Group Policies, on page 1619.

Step 9 Click **Next**.

Step 10 Select the **Secure Client Image** that the VPN users will use to connect to the remote access VPN.

The Secure Client provides secure SSL or IPSec (IKEv2) connections to the Secure Firewall Threat Defense device for remote users with full VPN profiling to corporate resources. After the remote access VPN policy is deployed on the Firewall Threat Defense device, VPN users can enter the IP address of the configured device interface in their browser to download and install the Secure Client.

For information about configuring the client profile and client modules, see Group Policy Secure Client Options, on page 1469.

Step 11 Click **Next**.

Step 12 Configure **Network Interface for Incoming VPN Access**.

Interface objects segment your network to help you manage and classify traffic flow. A security zone object simply groups interfaces. These groups may span multiple devices; you can also configure multiple zones interface objects on a single device. There are two types of interface objects:

- Security zones—An interface can belong to only one security zone.
- Interface groups—An interface can belong to multiple interface groups (and to one security zone).

(Optional) Check the **Enable DTLS on member interfaces** check box, if required. DTLS is applicable only for SSL protocol.

Step 13 Configure **Device Certificates**.

Device certificate (also called Identity certificate) identifies the VPN gateway to the remote access clients. Select a certificate which is used to authenticate the VPN gateway. From the **Certificate Enrollment** drop-down list, choose a certificate or click + to add a certificate to authenticate the VPN gateway.

Step 14 Configure **Service Access Control**.

You can use a service access object to control remote clients' access to VPN on Firewall Threat Defense devices with Version 7.7 or later. This object provides geolocation-based access control to clients before VPN authentication. By default, there is no access control for RAVPN, and remote clients can connect from any geolocation unless specified by a service access object. For more information, see Manage VPN Access of Remote Users Based on Geolocation, on page 1645and Configure a Service Access Object, on page 1372.

Step 15 Configure **Access Control for VPN Traffic** .

By default, all decrypted traffic in the VPN tunnel is subjected to the Access Control Policy. Check the **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** check box to bypass decrypted traffic from the Access Control Policy. This option bypasses the Access Control Policy inspection, but VPN filter ACL and authorization ACL downloaded from AAA server are still applied to VPN traffic.

Note
If you select this option, you need not update the access control policy for remote access VPN as specified in Update the Access Control Policy on the Secure Firewall Threat Defense Device, on page 1587.

Step 16 Click **Next**.

Step 17 View the **Summary** of the remote access VPN policy configuration.

The Summary page displays all the remote access VPN settings you have configured so far and provides links to the additional configurations that need to be performed before deploying the remote access VPN policy on the selected devices.

Click **Back** to make changes to the configuration, if required.

Step 18 Click **Finish** to complete the basic configuration for the remote access VPN policy.

When you complete the Remote Access VPN Policy Wizard, the policy listing page appears. Later, set up DNS configuration, configure access control for VPN users, and enable NAT exemption (if necessary) to complete a basic remote access VPN Policy configuration.

What to do next

Use the Remote Access VPN dashboard (**Overview > Dashboards > Remote Access VPN**) to monitor real-time data from active remote access VPN sessions on the devices. You can quickly determine problems related to user sessions and mitigate the problems for your network and users. For more information, see Remote Access VPN Dashboard, on page 1703.

Update the Access Control Policy on the Secure Firewall Threat Defense Device

Before deploying the remote access VPN policy, you must update the access control policy on the targeted Secure Firewall Threat Defense device with a rule that allows VPN traffic. The rule must allow all traffic coming in from the outside interface, with source as the defined VPN pool networks and destination as the corporate network.

Note If you have selected the **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** option on the Access Interface tab, you need not update the access control policy for remote access VPN.

Enable or disable the option for all your VPN connections. If you disable this option, make sure that the traffic is allowed by the access control policy or pre-filter policy.

For more information, see Configure Access Interfaces for Remote Access VPN, on page 1613.

Before you begin

Complete the remote access VPN policy configuration using the Remote Access VPN Policy wizard.

Procedure

Step 1 On your Secure Firewall Management Center web interface, choose **Policies** > **Access Control**.

Step 2 Click **Edit** on the access control policy that you want to update.

Step 3 Click **Add Rule** to add a new rule.

Step 4 Specify the **Name** for the rule and select **Enabled**.

Step 5 Select the **Action**, **Allow** or **Trust**.

Step 6 Select the following on the **Zones** tab:

a) Select the outside zone from Available Zones and click **Add to Source**.

b) Select the inside zone from Available Zones and click **Add to Destination**.

Step 7 Select the following on the **Networks** tab:

a) Select the inside network (inside interface and/or a corporate network) from Available networks and click **Add to Destination**.

b) Select the VPN address pool network from **Available Networks** and click **Add to Source Networks**.

Step 8 Configure other required access control rule settings and click **Add**.

Step 9 Save the rule and access control policy.

(Optional) Configure NAT Exemption

NAT exemption exempts addresses from translation and allows both translated and remote hosts to initiate connections with your protected hosts. Like identity NAT, you do not limit translation for a host on specific interfaces; you must use NAT exemption for connections through all interfaces. However, NAT exemption enables you to specify the real and destination addresses when determining the real addresses to translate (similar to policy NAT). Use static identity NAT to consider ports in the access list.

When you configure static identity NAT for remote access or site-to-site VPN, you must configure NAT with the route lookup option. Without route lookup, the Firewall Threat Defense sends traffic out of the interface specified in the NAT command, regardless of what the routing table says. For example, you do not want the Firewall Threat Defense to send the DHCP scope traffic through an incorrect interface; it will never return to the interface IP address. The route lookup option lets the Firewall Threat Defense send, or intercept, the traffic directly on the interface IP address instead of through the interface. For traffic from the VPN client to a host on the inside network, the route lookup option will still result in the correct egress interface (inside), so normal traffic flow is not affected.

Before you begin

Check if NAT is configured on the targeted devices where remote access VPN policy is deployed. If NAT is enabled on the targeted devices, you must define a NAT policy to exempt VPN traffic.

Procedure

Step 1 On your Secure Firewall Management Center web interface, click **Devices** > **NAT**.

Step 2 Select a NAT policy to update or click **New Policy** to create a NAT policy with a NAT rule to allow connections through all interfaces.

Step 3 Click **Add Rule** to add a NAT rule.

Step 4 On the Add NAT Rule window, select the following:

a) Select the NAT Rule as **Manual NAT Rule**.

b) Select the Type as **Static**.

c) Click **Interface Objects** and select the Source and destination interface objects.

Note
This interface object must be the same as the interface selected in the remote access VPN policy.

For more information, see Configure Access Interfaces for Remote Access VPN, on page 1613.

a) Click **Translation** and select the source and destination networks:

- **Original Source** and **Translated Source**
- **Original Destination** and **Translated Destination**

Step 5 On the Advanced tab, select **Do not proxy ARP on Destination Interface**.

Do not proxy ARP on Destination Interface—Disables proxy ARP for incoming packets to the mapped IP addresses. If you use addresses on the same network as the mapped interface, the system uses proxy ARP to answer any ARP requests for the mapped addresses, thus intercepting traffic destined for a mapped address. This solution simplifies routing because the device does not have to be the gateway for any additional networks. You can disable proxy ARP if desired, in which case you need to be sure to have proper routes on the upstream router.

Step 6 Click **OK**.

Configure DNS

Configure DNS on each Firewall Threat Defense device in order to use remote access VPN. Without DNS, the devices cannot resolve AAA server names, named URLs, and CA Servers with FQDN or Hostnames. It can only resolve IP addresses.

Procedure

Step 1 Configure DNS server details and domain-lookup interfaces using the Platform Settings. For more information, see DNS, on page 949 and DNS Server Group, on page 1384.

Step 2 Configure split-tunnel in group policy to allow DNS traffic through remote access VPN tunnel if the DNS server is reachable through VNP network. For more information, see Configure Group Policy Objects, on page 1466.

Add Secure Client Profile XML File

The Secure Client Profile is a group of configuration parameters stored in an XML file that the client uses to configure its operation and appearance. These parameters (XML tags) include the names and addresses of host computers and settings to enable more client features.

You can create the Secure Client Profile using the Secure Client Profile Editor, a GUI-based configuration tool that is available as part of the Secure Client software package. It is an independent program that you run outside of the Firewall Management Center. For more information about Secure Client Profile Editor, see Cisco Secure Client (including AnyConnect) Administrator Guide.

Before you begin

A Secure Firewall Threat Defense remote access VPN policy requires assignment of the Secure Client Profile to the VPN clients. You can attach the client profile to a group policy.

Download the Secure Client Profile Editor from Cisco Software Download Center.

Procedure

Step 1 Choose **Devices** > **Remote Access**.

Step 2 Click **Edit** on the remote access VPN policy which you want to update.

Step 3 Click **Edit** on the connection profile to which you want to add the Secure Client profile.

Step 4 Click **Edit Group Policy**. If you choose to add a new group policy, click **Add**.

Step 5 Choose **Secure Client** > **Profile**.

Step 6 Choose a profile from the **Client Profile** drop-down list. If you choose to add a new client profile, click **Add** and do the following:

a) Specify the profile **Name**.

b) Click **Browse** and select the Secure Client Profile XML file.

Note
For two-factor authentication, make sure that the timeout is set to 60 seconds or more in the Secure Client profile.

c) Click **Save**.

Step 7 Save your changes.

(Optional) Configure Split Tunneling

Split tunnel allows VPN connectivity to a remote network across a secure tunnel and also to a network outside VPN tunnel. Configure split tunneling if you want to allow your VPN users to access an outside network while they remain connected to the remote access VPN. To configure a split-tunnel list, you must create a Standard Access List or Extended Access List.

For more information, see Configuring Group Policies, on page 1619.

Procedure

Step 1 Choose **Devices** > **Remote Access**.

Step 2 Click **Edit** on the remote access VPN policy for which you want to configure split tunneling.

Step 3 Click **Edit** on the required connection profile.

Step 4 Click **Add** to add a group policy or click **Edit Group Policy**.

Step 5 Choose **General** > **Split Tunneling**.

Step 6 From the **IPv4 Split Tunneling** or **IPv6 Split Tunneling** list, select **Exclude networks specified below** and then select the networks that you want to exclude from VPN traffic.

The default setting allows all traffic over the VPN tunnel.

Step 7 Click **Standard Access List** or **Extended Access List**, and select an access list from the drop-down or add a new one.

Step 8 If you choose to add a new standard or extended access list, do the following:

a) Specify the **Name** for the new access list and click **Add**.
b) Choose **Allow** from the **Action** drop-down.
c) Select the network traffic that you want to allow over the VPN tunnel and click **Add**.

Step 9 Save your changes.

Related Topics

Access List, on page 1369

(Optional) Configure Dynamic Split Tunneling

Dynamic split tunneling allows you to fine-tune split tunneling based on DNS domain names. You can configure domains that must be included or excluded in the remote access VPN tunnel. Excluded domains are not blocked. Instead, traffic to those domains is kept outside the VPN tunnel. For example, you could send traffic to Cisco WebEx on the public Internet, thus freeing bandwidth in your VPN tunnel for traffic that is targeted to servers within your protected network. For more information about configuring this feature, see Configure AnyConnect Dynamic Split Tunnel on FTD Managed by FMC.

Before you begin

You can configure this feature using the Firewall Management Center and Firewall Threat Defense from versions 7.0 or later. If you have an older version of the Firewall Management Center, you can configure it using FlexConfig as instructed in the Advanced AnyConnect VPN Deployments for Firepower Threat Defense with FMC.

Procedure

Step 1 Configure the group policy to use Dynamic Split Tunnel.

a) Choose **Devices > Remote Access**.
b) Click **Edit** on the remote access VPN policy for which you want to configure dynamic split tunneling.
c) Click **Edit** on the required connection profile.
d) Click **Edit Group Policy**.

Step 2 Configure the Secure Client custom attribute in the **Add/Edit Group Policy** dialog box.

a) Click the Secure Client tab.
b) Click **Custom Attributes** and click **+**.
c) Choose **Dynamic Split Tunneling** from the **Secure Client Attribute** drop-down list.
d) Click **+** to create a new custom attribute object.
e) Enter the name for the custom attribute object.
f) **Include domains**—Specify domain names that will be included in the remote access VPN tunnel.

You can include domains in the tunnel that will be excluded based on IP addresses.

g) **Exclude domains**—Specify domain names that will be excluded from the remote access VPN.

Excluded domains are not blocked, traffic to these domains is kept outside the VPN tunnel.

h) Click **Save**.

i) Click **Add**.

Step 3 Verify the configured custom attribute and click **Save** to save the group policy.

Step 4 Click **Save** to save the connection profile.

Step 5 Click **Save** to save the remote access VPN policy.

What to do next

1. Deploy the configuration to Firewall Threat Defense.
2. Verify the configured dynamic split tunnel configuration on the Firewall Threat Defense and the Secure Client. For more information, see Verify Dynamic Split Tunneling Configuration, on page 1592.

Verify Dynamic Split Tunneling Configuration

On the Firewall Threat Defense

Use the following commands to verify the dynamic split tunneling configuration:

- **show running-config webvpn**
- **show running-config anyconnect-custom-data**
- **show running-config group-policy** <*group-policy-name*>

On the Secure Client

Click the Statistics ( ) icon and choose **VPN > Statistics**. You can confirm the domains under the Dynamic Split Exclusion/Inclusion category.

Verify the Configuration

Procedure

Step 1 Open a web browser on a machine on the outside network.

Step 2 Enter the URL of the Firewall Threat Defense remote access VPN gateway device.

Step 3 Enter the username and password when prompted, and click **Logon**.

Note
Connection to the VPN establishes automatically if you install Secure Client on the system.

The VPN prompts you to download Secure Client if Secure Client is not installed.

Step 4 Download Secure Client if it is not installed and connect to the VPN.
The Secure Client installs itself. On successful authentication, you establish connection to the Secure Firewall Threat Defense remote access VPN gateway. The remote access VPN enforces the applicable identity or QoS policy according to your VPN policy configuration.

Create a Copy of an Existing Remote Access VPN Policy

You can copy an existing remote access VPN policy to create a new one with all the settings, including the connection profiles and access interfaces. You can then assign devices to the new policy and deploy the VPN on the assigned devices as required.

Note Users with read-only permission for remote access VPN cannot create copy of the VPN. Users with read-only privileges in the domain can copy the remote access VPNs.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.
Step 2 Click **Copy** on the policy that you want to copy.
Step 3 Specify a **Name** for the new remote access VPN.
Step 4 Click **OK**.

What to do next

To assign devices to the new policy, see Set Target Devices for a Remote Access VPN Policy, on page 1593.

Set Target Devices for a Remote Access VPN Policy

After you create remote access VPN policy, you can assign the policy to threat defense devices.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.
Step 2 Click **Edit** (✎) next to the remote access VPN policy that you want to edit.
Step 3 Click **Policy Assignments**.
Step 4 Do any of the following:

- To assign a device, high-availability pair, or device group to the policy, select it in the **Available Devices** list and click **Add**. You can also drag and drop the available devices to select.

- To remove a device assignment, click **Delete** ( ) next to a device, high-availability pair, or device group in the **Selected Devices** list.

Step 5 Click **OK**.

Step 6 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Associate Local Realm with Remote Access VPN Policy

You can associate local realm to remote access VPN policy to enable local user authentication.

For information about creating and managing realms, see Manage a Realm, on page 2455.

For information about configuring local user authentication for remote access VPNs, see Configure AAA Settings for Remote Access VPN, on page 1597.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** ( ) next to the remote access VPN policy that you want to edit.

Step 3 Click the link next to **Local Realm**.

Step 4 Select the **Local Realm Server** from the list, or click **Add** to add a new local realm.

Step 5 Click **OK**.

Step 6 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Additional Remote Access VPN Configurations

Configure Connection Profile Settings

Remote Access VPN policy contains the connection profiles targeted for specific devices. These policies pertain to creating the tunnel itself, such as, how AAA is accomplished, and how addresses are assigned (DHCP or Address Pools) to VPN clients. They also include user attributes, which are identified in group policies configured on the Firewall Threat Defense device or obtained from a AAA server. A device also

provides a default connection profile named *DefaultWEBVPNGroup*. The connection profile that is configured using the wizard appears in the list.

If you decide to grant different rights to different groups of VPN users, then you can add specific connection profiles for each of the user groups and maintain multiple connection profiles in your remote access VPN policy.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select an existing remote access VPN policy in the list and click the corresponding **Edit** icon.

Step 3 Select a **Connection Profile** and click **Edit**.

Step 4 (Optional) If you choose to add new connection profile, click **Add**.

Step 5 Configure IP Addresses for VPN Clients.
Configure IP Addresses for VPN Clients, on page 1595

Step 6 (Optional) Update AAA Settings for remote access VPNs.
Configure AAA Settings for Remote Access VPN, on page 1597

Step 7 (Optional) Create or update Aliases.
Create or Update Aliases for a Connection Profile, on page 1612

Step 8 Save your changes.

Configure IP Addresses for VPN Clients

Client address assignment allows you to assign IP addresses for the remote access VPN users.

You can assign IP Address for remote VPN clients from the local IP address pools, DHCP Servers, and AAA servers. The AAA servers are assigned first, followed by others. Configure the **Client Address Assignment** policy in the **Advanced** tab to define the assignment criteria. The IP pools defined in this connection profile will only be used if no IP pools are defined in group policy associated with the connection profile, or the system default group policy **DfltGrpPolicy**.

IPv4 Address Pools—SSL VPN clients receive new IP addresses when they connect to the Firewall Threat Defense device. Address pools define a range of addresses that remote clients can receive. You can add a maximum of six pools for IPv4 and IPv6 addresses each.

Note You can use the IP address from the existing IP pools in the Firewall Management Center or create a new pool using the **Add** option. Also, you can create an IP pool in Firewall Management Center using the **Objects** > **Object Management** > **Address Pools** path. For more information, see Address Pools, on page 1374.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.
Existing remote access policies are listed.

Step 2 Select a remote access VPN policy and click the edit icon.

Step 3 Select the connection profile that you want to update and click the edit icon.

Step 4 Under the **Client Address Assignment** tab, do the following:

Step 5 Click + next to **Address Pools**:

a) Click + next to **Address Pools** to add IP addresses, and select **IPv4** or **IPv6** to add the corresponding address pool. Select the IP address pool from **Available Pools** and click **Add**.

Note
If you share your remote access VPN policy among multiple Secure Firewall Threat Defense devices, bear in mind that all devices share the same address pool unless you use device-level object overrides to replace the global definition with a unique address pool for each device. Unique address pools are required to avoid overlapping addresses in cases where the devices are not using NAT.

b) Click + next to **Available Pools** in the **Address Pools** window to add a new IPv4 or IPv6 address pool. When you choose the IPv4 pool, provide a starting and ending IP address. When you choose to include a new IPv6 address pool, enter **`Number of Addresses`** in the range 1-16384. Select the **Allow Overrides** option to avoid conflicts with IP address when objects are shared across many devices. For more information, see Address Pools, on page 1374.

c) Click **OK**.

If you plan to edit the IP address pools, we recommend that you perform the following steps during a maintenance window:

1. Unassign the device from the remote access VPN.
2. Select the device and click **Deploy**.

   This deployment removes all the remote access VPN configurations from the device, terminates the remote access VPN sessions, the sessions are not reestablished.
3. Click the edit icon next to the IP address pools to edit it, edit any other remote access VPN configurations, if required, on the Firewall Management Center.
4. Assign the device to the updated remote access VPN policy.
5. Deploy the configurations on the device.

   The remote access VPN clients can connect to the device after the maintenance window.

Step 6 Click + next to **DHCP Servers** to add DHCP servers:

Note
The DHCP server address can be configured only with IPv4 address.

a) Specify the name and DHCP (Dynamic Host Configuration Protocol) server address as network objects. Click **Add** to choose the server from the object list. Click **Delete** to delete a DHCP server.

b) Click **Add** in the **New Objects** page to add a new network object. Enter the new object name, description, network, and select the **Allow Overrides** option as applicable. For more information, see Creating Network Objects, on page 1403 and Allowing Object Overrides, on page 1362.

c) Click **OK**.

Step 7 Click **Save**.

Related Topics

Configure Connection Profile Settings, on page 1594

Configure AAA Settings for Remote Access VPN

Before you begin

- Ensure that the required machine and user certificates are deployed on the endpoints. For information about Firewall Threat Defense certificates, see Managing Firewall Threat Defense Certificates, on page 1486.
- Configure Secure Client profiles with required certificates. For more information, see *Cisco Secure Client (including AnyConnect) Administrator Guide*.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select an existing remote access VPN policy in the list and click the corresponding **Edit** icon.

Step 3 Select a connection profile to update AAA settings, click **Edit** > **AAA**.

Step 4 Select the following for **Authentication**:

- **Authentication Method**—Determines how a user is identified before being allowed access to the network and network services. It controls access by requiring valid user credentials, which are typically a username and password. It may also include the certificate from the client. Supported authentication methods are **AAA only**, **Client Certificate only**, and **AAA + Client Certificate**.

  When you select the **Authentication Method** as:

  - **AAA Only**—If you select the **Authentication Server** as **RADIUS**, by default, the Authorization Server has the same value. Select the **Accounting Server** from the drop-down list. Whenever you select **AD** and **LDAP** from the Authentication Server drop-down list, you must select the **Authorization Server** and **Accounting Server** manually.
  - **SAML**—Each user is authenticated using the SAML single sign-on server. For more information, see Single Sign-On Authentication with SAML 2.0, on page 1666.

    Override Identity Provider Certificate—Select to override the primary identity provider certificate of the SAML provider with an IdP certificate specific to a connection profile or SAML application. Select the IdP certificate from the drop-down.

    Microsoft Azure can support multiple applications for the same entity ID. Each application (mapped to a different connection profile) requires a unique certificate. If you want to retain an existing entity ID for the single-sign-on object in current connection profile and use a different IdP certificate, you can select this option.

    This enables support for multiple SAML applications per Microsoft Azure SAML identity provider.

    The primary identity certificate is configured in the single sign-on server object.

    For information about configuring a single sign-on server object, see Add a Single Sign-on Server, on page 1367.

    Choose your **SAML Login Experience** to configure a browser for SAML web authentication:

    - **VPN client embedded browser**—Choose this option to use the browser embedded with the VPN client for web authentication. The authentication applies to the VPN connection only.

- **Default OS Browser**—Choose this option to configure the operating system that default or native browser that supports WebAuthN (FIDO2 standard for web authentication). This option enables single sign-on(SSO) support for web authentication methods such as biometric authentication.

  The default browser requires an external browser package for web authentication. The package Default-External-Browser-Package is configured by default. You can change the default external browser package by editing a remote access VPN policy and selecting the file under **Advanced** > **Secure Client Images** > **Package File**.

  You can also add a new package file by selecting. **Objects** > **Object Management** > **VPN** > **Secure Client File** > **Add Secure Client File**.

- **Client Certificate Only**—Each user is authenticated with a client certificate. The client certificate must be configured on VPN client endpoints. By default, the user name is derived from the client certificate fields CN and OU. If the user name is specified in other fields in the client certificate, use 'Primary' and 'Secondary' field to map appropriate fields.

  Select **Enable multiple certificate authentication** to authenticate the VPN client using the machine and user certificates.

  If have enabled multiple certificate authentication, you can select one of the following certificates to map the username and authenticate the VPN user:

  - **First Certificate**—Select this option to map the username from the machine certificate sent from the VPN client.
  - **Second Certificate**—Select this option to map the username from the user certificate sent from the client.

  Note
  If you do not enable multiple certificate authentication, the user certificate (second certificate) is used for authentication by default.

  If you select the **Map specific field** option, which includes the username from the client certificate, the **Primary** and **Secondary** fields display default values: **CN (Common Name)** and **OU (Organisational Unit)** respectively. If you select the **Use entire DN as username** option, the system automatically retrieves the user identity. A distinguished name (DN) is a unique identification, made up of individual fields used as the identifier when matching users to a connection profile. DN rules are used for enhanced certificate authentication.

  The primary and Secondary fields pertaining to the **Map specific field** option contain these common values:

  - C (Country)
  - CN (Common Name)
  - DNQ (DN Qualifier)
  - EA (Email Address)
  - GENQ (Generational Qualifier)
  - GN (Given Name)
  - I (Initial)

- L (Locality)
- N (Name)
- O (Organisation)
- OU (Organisational Unit)
- SER (Serial Number)
- SN (Surname)
- SP (State Province)
- T (Title)
- UID (User ID)
- UPN (User Principal Name)

- **Client Certificate & AAA**— Each user is authenticated with both a client certificate and AAA server. Select the required certificate and AAA configurations for authentication.

  Whichever authentication method you choose, select or deselect **Allow connection only if user exists in authorization database**.

- **Client Certificate & SAML**— Each user is authenticated with both a client certificate and SAML server. Select the required certificate and SAML configurations for authentication.
  - **Allow connection only if username from certificate and SAML are the same**—Select to allow VPN connection only if the username from the certificate matches the SAML single sign-on username.
  - **Use username from client certificate for Authorization**—When you choose the option to pick username from client certificate for authorization, you must configure the fields to pick from the client certificate.

    You can choose to map a specific filed as the username or use the entire distinguished name (DN) for authorization:
    - **Map specific field**— Select to include the username from the client certificate; the **Primary** and **Secondary** fields display default values: **CN (Common Name)** and **OU (Organisational Unit)** respectively.
    - **Use entire DN as username**— The system automatically retrieves the user identity for authorization.

  You can also create a Dynamic Access Policy (DAP) to match user-specific SAML assertion attributes or username to DAP certificate attributes. See Configure AAA Criteria Settings for DAP, on page 1693.

- **Authentication Server**—Authentication is the way a user is identified before being allowed access to the network and network services. Authentication requires valid user credentials, a certificate, or both. You can use authentication alone, or with authorization and accounting.

  Select an authentication server from the list if you have added a server already or else create an authentication server:

- **LOCAL**—Use a local database from the Firewall Threat Defense for user authentication. To configure local authentication, Firewall Threat Defense must be Version 7.0 and later.
  - **Local Realm**—Select a local realm or click **Add** to configure a realm. See Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.
- **Realm**—Configure an LDAP or AD realm. See Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.
- **RADIUS Server Group**—Add a RADIUS server group object with RADIUS servers. See Add a RADIUS Server Group, on page 1364.
- **Single Sign-On Server**—Create a single sign-on server object for SAML authentication. See Add a Single Sign-on Server, on page 1367.

Fallback to LOCAL Authentication— The user is authenticated using the local database and the VPN tunnel can be established even if the AAA server group is unavailable, provided that the local database is configured.

- **Use secondary authentication** — Secondary authentication is configured in addition to primary authentication to provide additional security for VPN sessions. Secondary authentication is applicable only to **AAA only** and **Client Certificate & AAA** authentication methods.

  Secondary authentication is an optional feature that requires a VPN user to enter two sets of username and password on the Secure Client login screen. You can also configure to pre-fill the secondary username from the authentication server or client certificate. Remote access VPN authentication is granted only if both primary and secondary authentications are successful. VPN authentication is denied if any one of the authentication servers is not reachable or one authentication fails.

  You must configure a secondary authentication server group (AAA server) for the second username and password before configuring secondary authentication. For example, you can set the primary authentication server to an LDAP or Active Directory realm and the secondary authentication to a RADIUS server.

  Note
  By default, secondary authentication is not required.

  Authentication Server— Secondary authentication server to provide secondary username and password for VPN users.

  - **Fallback to LOCAL Authentication**— This user is authenticated using the local database and the VPN tunnel can be established even if the AAA server group is unavailable, provided that the local database is configured.

  Select the following under **Username for secondary authentication**:

  - **Prompt**: Prompts the users to enter the username and password while logging on to VPN gateway.
  - **Use primary authentication username**: The username is taken from the primary authentication server for both primary and secondary authentication; you must enter two passwords.
  - **Map username from client certificate**: Prefills the secondary username from the client certificate.

    If you have enabled multiple certificate authentication, you can select one of the following certificates:

    - **First Certificate**— Select this option to map the username from the machine certificate sent from the VPN client.

- **Second Certificate**— Select this option to map the username from the user certificate sent from the client.
- If you select **Map specific field** option, which includes the username from the client certificate. The **Primary** and **Secondary** fields display default values: **CN (Common Name)** and **OU (Organisational Unit)** respectively. If you select the **Use entire DN (Distinguished Name) as username** option, the system automatically retrieves the user identity.

  See **Authentication Method** descriptions for more information about primary and secondary field mapping.
- **Prefill username from certificate on user login window**: Prefills the secondary username from the client certificate when the user connects via Secure Client.
  - **Hide username in login window**: The secondary username is pre-filled from the client certificate, but hidden to the user so that the user does not modify the pre-filled username.

- **Use secondary username for VPN session**: The secondary username is used for reporting user activity during a VPN session.

Step 5 Select the following for **Authorization**:

- **Authorization Server**—After authentication is complete, authorization controls the services and commands available to each authenticated user. Authorization works by assembling a set of attributes that describe what the user is authorized to perform, their actual capabilities, and restrictions. When you do not use authorization, authentication alone provides the same access to all authenticated users. Authorization requires authentication.

  To know more about how remote access VPN authorization works, see Understanding Policy Enforcement of Permissions and Attributes, on page 1578.

  When a RADIUS Server is configured for user authorization in the connection profile, the remote access VPN system administrator can configure multiple authorization attributes for users or user-groups. The authorization attributes that are configured on the RADIUS server can be specific for a user or a user-group. Once users are authenticated, these specific authorization attributes are pushed to the Firewall Threat Defense device.

  Note
  The AAA server attributes obtained from the authorization server override the attribute values that may have been previously configured on the group policy or the connection profile.

- Check **Allow connection only if user exists in authorization database** if desired.

  When enabled, the system checks the username of the client must exist in the authorization database to allow a successful connection. If the username does not exist in the authorization database, then the connection is denied.

- When you select a realm as the Authorization Server, you must configure an LDAP attribute map. You can choose a single server for authentication and authorization or a different servers. Click **Configure LDAP Attribute Map** to add LDAP attribute maps for authorization.

  Note
  Firewall Threat Defense does not support SAML Identity Provider as the Authorization server. If the Active Directory behind the SAML Identity Provider is reachable via Firewall Management Center and Firewall Threat Defense, you can configure authorization following these steps:

- Add realm for the AD Server. See Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.
- Select the realm object as the Authorization Server in remote access VPN connection profile.
- Configure LDAP attribute map for the selected realm.

For information about configuring LDAP attribute maps, see Configuring LDAP Attribute Mapping, on page 1619.

Step 6 Select the following for **Accounting**:

- **Accounting Server**—Accounting is used to track the services that users are accessing and the amount of network resources they are consuming. When AAA accounting is activated, the network access server reports user activity to the RADIUS server. Accounting information includes when sessions start and stop, usernames, the number of bytes that pass through the device for each session, the services used, and the duration of each session. This data can then be analyzed for network management, client billing, or auditing. You can use accounting alone or together with authentication and authorization.

  Specify the RADIUS Server Group object that will be used to account for the remote access VPN session.

Step 7 Select the following **Advanced Settings**:

- **Strip Realm from username**—Select to remove the realm from the username before passing the username on to the AAA server. For example, if you select this option and provide *domain\username*, the domain is stripped off from the username and sent to AAA server for authentication. By default this option is unchecked.
- **Strip Group from username**—Select to remove the group name from the username before passing the username on to the AAA server. By default this option is unchecked.

  Note
  A realm is an administrative domain. Enabling these options allows the authentication to be based on the username alone. You can enable any combination of these options. However, you must select both check boxes if your server cannot parse delimiters.
- **Password Management**—Enable managing the password for the remote access VPN users. Select to notify ahead of the password expiry or on the day the password expires.

Step 8 Click **Save**.

Related Topics

Understanding Policy Enforcement of Permissions and Attributes, on page 1578
Manage a Realm, on page 2455

RADIUS Server Attributes for Secure Firewall Threat Defense

The Firewall Threat Defense device supports applying user authorization attributes (also called user entitlements or permissions) to VPN connections from the external RADIUS server that are configured for authentication and/or authorization in the remote access VPN policy.

Note Secure Firewall Threat Defense devices support attributes with vendor ID 3076.

The following user authorization attributes are sent to the Firewall Threat Defense device from the RADIUS server.

- RADIUS attributes 146 and 150 are sent from Firewall Threat Defense devices to the RADIUS server for authentication and authorization requests.
- All three (146, 150, and 151) attributes are sent from Firewall Threat Defense devices to the RADIUS server for accounting start, interim-update, and stop requests.

Table 81: RADIUS Attributes Sent from Secure Firewall Threat Defense to RADIUS Server

Attribute	Attribute Number	Syntax, Type	Single or Multi-valued	Description or Value
Connection Profile Name or Tunnel Group Name	146	String	Single	1-253 characters
Client Type	150	Integer	Single	2 = Secure Client SSL VPN, 6 = Secure Client IPsec VPN (IKEv2)
Session Type	151	Integer	Single	1 = Secure Client SSL VPN, 2 = Secure Client IPsec VPN (IKEv2)

Table 82: Supported RADIUS Authorization Attributes

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
Access-Hours	Y	1	String	Single	Name of the time range, for example, Busine
Access-List-Inbound	Y	86	String	Single	Both of the Access-List attributes take the na ACL that is configured on the Firewall Threa device. Create these ACLs using the Smart CL Access List object type (select **Device** > **Adv Configuration** > **Smart CLI** > **Objects**). These ACLs control traffic flow in the inbou entering the Firewall Threat Defense device) (traffic leaving the Firewall Threat Defense direction.
Access-List-Outbound	Y	87	String	Single	
Address-Pools	Y	217	String	Single	The name of a network object defined on the Threat Defense device that identifies a subnet be used as the address pool for clients conne remote access VPN. Define the network obje **Objects** page and then associate the network a group policy or a connection profile.

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
Allow-Network-Extension-Mode	Y	64	Boolean	Single	0 = Disabled 1 = Enabled
Authenticated-User-Idle-Timeout	Y	50	Integer	Single	1-35791394 minutes
Authorization-DN-Field	Y	67	String	Single	Possible values: UID, OU, O, CN, L, SP, C, EA, GN, SN, I, GENQ, DNQ, SER, use-entire-name
Authorization-Required		66	Integer	Single	0 = No 1 = Yes
Authorization-Type	Y	65	Integer	Single	0 = None 1 = RADIUS 2 = LDAP
Banner1	Y	15	String	Single	Banner string to display for Cisco VPN remote a sessions: IPsec IKEv1, Secure Client SSL-TLS/DTLS/IKEv2, and Clientless SSL
Banner2	Y	36	String	Single	Banner string to display for Cisco VPN remote a sessions: IPsec IKEv1, Secure Client SSL-TLS/DTLS/IKEv2, and Clientless SSL. The string is concatenated to the Banner1 string , if con
Cisco-IP-Phone-Bypass	Y	51	Integer	Single	0 = Disabled 1 = Enabled
Cisco-LEAP-Bypass	Y	75	Integer	Single	0 = Disabled 1 = Enabled
Client Type	Y	150	Integer	Single	1 = Cisco VPN Client (IKEv1) 2 = Secure Clien VPN 3 = Clientless SSL VPN 4 = Cut-Through- = L2TP/IPsec SSL VPN 6 = Secure Client IPsec (IKEv2)
Client-Type-Version-Limiting	Y	77	String	Single	IPsec VPN version number string
DHCP-Network-Scope	Y	61	String	Single	IP Address
Extended-Authentication-On-Rekey	Y	122	Integer	Single	0 = Disabled 1 = Enabled
Framed-Interface-Id	Y	96	String	Single	Assigned IPv6 interface ID. Combines with Framed-IPv6-Prefix to create a complete assigne address. For example: Framed-Interface-ID=1:1 combined with Framed-IPv6-Prefix=2001:0db8::/ the assigned IP address 2001:0db8::1:1:1:1.
Framed-IPv6-Prefix	Y	97	String	Single	Assigned IPv6 prefix and length. Combines with Framed-Interface-Id to create a complete assigne address. For example: prefix 2001:0db8::/64 con with Framed-Interface-Id=1:1:1:1 gives the IP a 2001:0db8::1:1:1:1. You can use this attribute to an IP address without using Framed-Interface-Id assigning the full IPv6 address with prefix lengt for example, Framed-IPv6-Prefix=2001:0db8::1

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
Group-Policy	Y	25	String	Single	Sets the group policy for the remote access V You can use one of the following formats: • *group policy name* • OU=*group policy name* • OU=*group policy name*;
IE-Proxy-Bypass-Local		83	Integer	Single	0 = None 1 = Local
IE-Proxy-Exception-List		82	String	Single	New line (\n) separated list of DNS domains
IE-Proxy-PAC-URL	Y	133	String	Single	PAC address string
IE-Proxy-Server		80	String	Single	IP address
IE-Proxy-Server-Policy		81	Integer	Single	1 = No Modify 2 = No Proxy 3 = Auto detec Concentrator Setting
IKE-KeepAlive-Confidence-Interval	Y	68	Integer	Single	10-300 seconds
IKE-Keepalive-Retry-Interval	Y	84	Integer	Single	2-10 seconds
IKE-Keep-Alives	Y	41	Boolean	Single	0 = Disabled 1 = Enabled
Intercept-DHCP-Configure-Msg	Y	62	Boolean	Single	0 = Disabled 1 = Enabled
IPsec-Allow-Passwd-Store	Y	16	Boolean	Single	0 = Disabled 1 = Enabled
IPsec-Authentication		13	Integer	Single	0 = None 1 = RADIUS 2 = LDAP (authoriza = NT Domain 4 = SDI 5 = Internal 6 = RAD Expiry 7 = Kerberos/Active Directory
IPsec-Auth-On-Rekey	Y	42	Boolean	Single	0 = Disabled 1 = Enabled
IPsec-Backup-Server-List	Y	60	String	Single	Server Addresses (space delimited)
IPsec-Backup-Servers	Y	59	String	Single	1 = Use Client-Configured list 2 = Disable and list 3 = Use Backup Server list
IPsec-Client-Firewall-Filter-Name		57	String	Single	Specifies the name of the filter to be pushed as firewall policy
IPsec-Client-Firewall-Filter-Optional	Y	58	Integer	Single	0 = Required 1 = Optional
IPsec-Default-Domain	Y	28	String	Single	Specifies the single default domain name to client (1-255 characters).
IPsec-IKE-Peer-ID-Check	Y	40	Integer	Single	1 = Required 2 = If supported by peer certifi not check

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
IPsec-IP-Compression	Y	39	Integer	Single	0 = Disabled 1 = Enabled
IPsec-Mode-Config	Y	31	Boolean	Single	0 = Disabled 1 = Enabled
IPsec-Over-UDP	Y	34	Boolean	Single	0 = Disabled 1 = Enabled
IPsec-Over-UDP-Port	Y	35	Integer	Single	4001- 49151. The default is 10000.
IPsec-Required-Client-Firewall-Capability	Y	56	Integer	Single	0 = None 1 = Policy defined by remote FW Are-You-There (AYT) 2 = Policy pushed CPP 4 = from server
IPsec-Sec-Association		12	String	Single	Name of the security association
IPsec-Split-DNS-Names	Y	29	String	Single	Specifies the list of secondary domain names to the client (1-255 characters).
IPsec-Split-Tunneling-Policy	Y	55	Integer	Single	0 = No split tunneling 1 = Split tunneling 2 = Loc permitted
IPsec-Split-Tunnel-List	Y	27	String	Single	Specifies the name of the network or ACL that d the split tunnel inclusion list.
IPsec-Tunnel-Type	Y	30	Integer	Single	1 = LAN-to-LAN 2 = Remote access
IPsec-User-Group-Lock		33	Boolean	Single	0 = Disabled 1 = Enabled
IPv6-Address-Pools	Y	218	String	Single	Name of IP local pool-IPv6
IPv6-VPN-Filter	Y	219	String	Single	ACL value
L2TP-Encryption		21	Integer	Single	Bitmap: 1 = Encryption required 2 = 40 bits 4 = 8 = Stateless-Req 15= 40/128-Encr/Stateless-Re
L2TP-MPPC-Compression		38	Integer	Single	0 = Disabled 1 = Enabled
Member-Of	Y	145	String	Single	Comma-delimited string, for example: `Engineering, Sales` An administrative attribute that can be used in d access policies. It does not set a group policy.
MS-Client-Subnet-Mask	Y	63	Boolean	Single	An IP address
NAC-Default-ACL		92	String		ACL
NAC-Enable		89	Integer	Single	0 = No 1 = Yes
NAC-Revalidation-Timer		91	Integer	Single	300-86400 seconds
NAC-Settings	Y	141	String	Single	Name of the NAC policy

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
NAC-Status-Query-Timer		90	Integer	Single	30-1800 seconds
Perfect-Forward-Secrecy-Enable	Y	88	Boolean	Single	0 = No 1 = Yes
PPTP-Encryption		20	Integer	Single	Bitmap: 1 = Encryption required 2 = 40 bits 8 = Stateless-Required 15= 40/128-Encr/Sta
PPTP-MPPC-Compression		37	Integer	Single	0 = Disabled 1 = Enabled
Primary-DNS	Y	5	String	Single	An IP address
Primary-WINS	Y	7	String	Single	An IP address
Privilege-Level	Y	220	Integer	Single	An integer between 0 and 15.
Required-Client- Firewall-Vendor-Code	Y	45	Integer	Single	1 = Cisco Systems (with Cisco Integrated Cl Zone Labs 3 = NetworkICE 4 = Sygate 5 = Cis (with Cisco Intrusion Prevention Security Ag
Required-Client-Firewall-Description	Y	47	String	Single	String
Required-Client-Firewall-Product-Code	Y	46	Integer	Single	Cisco Systems Products: 1 = Cisco Intrusion Prevention Security Age Integrated Client (CIC) Zone Labs Products: 1 = Zone Alarm 2 = Zon 3 = Zone Labs Integrity NetworkICE Product: 1 = BlackIce Defende Sygate Products: 1 = Personal Firewall 2 = P Firewall Pro 3 = Security Agent
Required-Individual-User-Auth	Y	49	Integer	Single	0 = Disabled 1 = Enabled
Require-HW-Client-Auth	Y	48	Boolean	Single	0 = Disabled 1 = Enabled
Secondary-DNS	Y	6	String	Single	An IP address
Secondary-WINS	Y	8	String	Single	An IP address
SEP-Card-Assignment		9	Integer	Single	Not used
Session Subtype	Y	152	Integer	Single	0 = None 1 = Clientless 2 = Client 3 = Clien Session Subtype applies only when the Sessi (151) attribute has the following values: 1, 2

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
Session Type	Y	151	Integer	Single	0 = None 1 = Secure Client SSL VPN 2 = Secur IPSec VPN (IKEv2) 3 = Clientless SSL VPN 4 = Clientless Email Proxy 5 = Cisco VPN Client (I = IKEv1 LAN-LAN 7 = IKEv2 LAN-LAN 8 = VI Balancing
Simultaneous-Logins	Y	2	Integer	Single	0-2147483647
Smart-Tunnel	Y	136	String	Single	Name of a Smart Tunnel
Smart-Tunnel-Auto	Y	138	Integer	Single	0 = Disabled 1 = Enabled 2 = AutoStart
Smart-Tunnel-Auto-Signon-Enable	Y	139	String	Single	Name of a Smart Tunnel Auto Signon list appen the domain name
Strip-Realm	Y	135	Boolean	Single	0 = Disabled 1 = Enabled
SVC-Ask	Y	131	String	Single	0 = Disabled 1 = Enabled 3 = Enable default ser Enable default clientless (2 and 4 not used)
SVC-Ask-Timeout	Y	132	Integer	Single	5-120 seconds
SVC-DPD-Interval-Client	Y	108	Integer	Single	0 = Off 5-3600 seconds
SVC-DPD-Interval-Gateway	Y	109	Integer	Single	0 = Off) 5-3600 seconds
SVC-DTLS	Y	123	Integer	Single	0 = False 1 = True
SVC-Keepalive	Y	107	Integer	Single	0 = Off 15-600 seconds
SVC-Modules	Y	127	String	Single	String (name of a module)
SVC-MTU	Y	125	Integer	Single	MTU value 256-1406 in bytes
SVC-Profiles	Y	128	String	Single	String (name of a profile)
SVC-Rekey-Time	Y	110	Integer	Single	0 = Disabled 1-10080 minutes
Tunnel Group Name	Y	146	String	Single	1-253 characters
Tunnel-Group-Lock	Y	85	String	Single	Name of the tunnel group or “none”
Tunneling-Protocols	Y	11	Integer	Single	1 = PPTP 2 = L2TP 4 = IPSec (IKEv1) 8 = L2T 16 = WebVPN 32 = SVC 64 = IPsec (IKEv2) 8 a mutually exclusive. 0 - 11, 16 - 27, 32 - 43, 48 - legal values.
Use-Client-Address		17	Boolean	Single	0 = Disabled 1 = Enabled
VLAN	Y	140	Integer	Single	0-4094
WebVPN-Access-List	Y	73	String	Single	Access-List name

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
WebVPN ACL	Y	73	String	Single	Name of a WebVPN ACL on the device
WebVPN-ActiveX-Relay	Y	137	Integer	Single	0 = Disabled Otherwise = Enabled
WebVPN-Apply-ACL	Y	102	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Auto-HTTP-Signon	Y	124	String	Single	Reserved
WebVPN-Citrix-Metaframe-Enable	Y	101	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Content-Filter-Parameters	Y	69	Integer	Single	1 = Java ActiveX 2 = Java Script 4 = Image 8 in images
WebVPN-Customization	Y	113	String	Single	Name of the customization
WebVPN-Default-Homepage	Y	76	String	Single	A URL such as http://example-example.com
WebVPN-Deny-Message	Y	116	String	Single	Valid string (up to 500 characters)
WebVPN-Download_Max-Size	Y	157	Integer	Single	0x7fffffff
WebVPN-File-Access-Enable	Y	94	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-File-Server-Browsing-Enable	Y	96	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-File-Server-Entry-Enable	Y	95	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Group-based-HTTP/HTTPS-Proxy-Exception-List	Y	78	String	Single	Comma-separated DNS/IP with an optional v (for example *.cisco.com, 192.168.1.*, wwwin
WebVPN-Hidden-Shares	Y	126	Integer	Single	0 = None 1 = Visible
WebVPN-Home-Page-Use-Smart-Tunnel	Y	228	Boolean	Single	Enabled if clientless home page is to be rende Smart Tunnel.
WebVPN-HTML-Filter	Y	69	Bitmap	Single	1 = Java ActiveX 2 = Scripts 4 = Image 8 =
WebVPN-HTTP-Compression	Y	120	Integer	Single	0 = Off 1 = Deflate Compression
WebVPN-HTTP-Proxy-IP-Address	Y	74	String	Single	Comma-separated DNS/IP:port, with http= c prefix (for example http=10.10.10.10:80, https=11.11.11.11:443)
WebVPN-Idle-Timeout-Alert-Interval	Y	148	Integer	Single	0-30. 0 = Disabled.
WebVPN-Keepalive-Ignore	Y	121	Integer	Single	0-900
WebVPN-Macro-Substitution	Y	223	String	Single	Unbounded.
WebVPN-Macro-Substitution	Y	224	String	Single	Unbounded.
WebVPN-Port-Forwarding-Enable	Y	97	Integer	Single	0 = Disabled 1 = Enabled

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
WebVPN-Port-Forwarding-Exchange-Proxy-Enable	Y	98	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Port-Forwarding-HTTP-Proxy	Y	99	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-Port-Forwarding-List	Y	72	String	Single	Port forwarding list name
WebVPN-Port-Forwarding-Name	Y	79	String	Single	String name (example, “Corporate-Apps”). This text replaces the default string, “Application A on the clientless portal home page.
WebVPN-Post-Max-Size	Y	159	Integer	Single	0x7fffffff
WebVPN-Session-Timeout-Alert-Interval	Y	149	Integer	Single	0-30. 0 = Disabled.
WebVPN Smart-Card-Removal-Disconnect	Y	225	Boolean	Single	0 = Disabled 1 = Enabled
WebVPN-Smart-Tunnel	Y	136	String	Single	Name of a Smart Tunnel
WebVPN-Smart-Tunnel-Auto-Sign-On	Y	139	String	Single	Name of a Smart Tunnel auto sign-on list appen the domain name
WebVPN-Smart-Tunnel-Auto-Start	Y	138	Integer	Single	0 = Disabled 1 = Enabled 2 = Auto Start
WebVPN-Smart-Tunnel-Tunnel-Policy	Y	227	String	Single	One of “e networkname,” “i networkname,” or “a networkname is the name of a Smart Tunnel netw e indicates the tunnel excluded, i indicates the tu specified, and a indicates all tunnels.
WebVPN-SSL-VPN-Client-Enable	Y	103	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SSL-VPN-Client-Keep- Installation	Y	105	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SSL-VPN-Client-Required	Y	104	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SSO-Server-Name	Y	114	String	Single	Valid string
WebVPN-Storage-Key	Y	162	String	Single	
WebVPN-Storage-Objects	Y	161	String	Single	
WebVPN-SVC-Keepalive-Frequency	Y	107	Integer	Single	15-600 seconds, 0=Off
WebVPN-SVC-Client-DPD-Frequency	Y	108	Integer	Single	5-3600 seconds, 0=Off
WebVPN-SVC-DTLS-Enable	Y	123	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-SVC-DTLS-MTU	Y	125	Integer	Single	MTU value is from 256-1406 bytes.
WebVPN-SVC-Gateway-DPD-Frequency	Y	109	Integer	Single	5-3600 seconds, 0=Off
WebVPN-SVC-Rekey-Time	Y	110	Integer	Single	4-10080 minutes, 0=Off

Attribute Name	Firewall Threat Defense	Attr. No.	Syntax/Type	Single or Multi- Valued	Description or Value
WebVPN-SVC-Rekey-Method	Y	111	Integer	Single	0 (Off), 1 (SSL), 2 (New Tunnel)
WebVPN-SVC-Compression	Y	112	Integer	Single	0 (Off), 1 (Deflate Compression)
WebVPN-UNIX-Group-ID (GID)	Y	222	Integer	Single	Valid UNIX group IDs
WebVPN-UNIX-User-ID (UIDs)	Y	221	Integer	Single	Valid UNIX user IDs
WebVPN-Upload-Max-Size	Y	158	Integer	Single	0x7fffffff
WebVPN-URL-Entry-Enable	Y	93	Integer	Single	0 = Disabled 1 = Enabled
WebVPN-URL-List	Y	71	String	Single	URL list name
WebVPN-User-Storage	Y	160	String	Single	
WebVPN-VDI	Y	163	String	Single	List of settings

Table 83: RADIUS Attributes Sent to Secure Firewall Threat Defense

Attribute	Attribute Number	Syntax, Type	Single or Multi-valued	Description or Value
Address-Pools	217	String	Single	The name of a network object defined on the Firewall Threat Defense device that identifies a subnet, which will be used as the address pool for clients connecting to the remote access VPN. Define the network object on the **Objects** page.
Banner1	15	String	Single	The banner to display when the user logs in.
Banner2	36	String	Single	The second part of the banner to display when the user logs in. Banner2 is appended to Banner1.
Downloadable ACLs	Cisco-AV-Pair	merge-dacl {before-avpair | after-avpair}		Supported via Cisco-AV-Pair configuration.
Filter ACLs	86, 87	String	Single	Filter ACLs are referred to by ACL name in the RADIUS server. It requires the ACL configuration to be already present on the Firewall Threat Defense device, so that it can be used during RADIUS authorization. 86=Access-List-Inbound 87=Access-List-Outbound

Attribute	Attribute Number	Syntax, Type	Single or Multi-valued	Description or Value
Group-Policy	25	String	Single	The group policy to use in the connection. You must create the group policy on the remote access VPN **Group Policy** page. You can use one of the following formats: • *group policy name* • OU=*group policy name* • OU=*group policy name*;
Simultaneous-Logins	2	Integer	Single	The number of separate simultaneous connections the user is allowed to establish, 0 - 2147483647.
VLAN	140	Integer	Single	The VLAN on which to confine the user's connection, 0 - 4094. You must also configure this VLAN on a subinterface on the Firewall Threat Defense device.

You must set the values of the IE-Proxy-Server-Method attribute returned from ISE to one of the following:

- IE_PROXY_METHOD_PACFILE: 8
- IE_PROXY_METHOD_PACFILE_AND_AUTODETECT: 11
- IE_PROXY_METHOD_PACFILE_AND_USE_SERVER: 12
- IE_PROXY_METHOD_PACFILE_AND_AUTODETECT_AND_USE_SERVER: 15

Firewall Threat Defense will deliver a proxy setting only if one of the above values is used for the IE-Proxy-Server-Method attribute.

Create or Update Aliases for a Connection Profile

Aliases contain alternate names or URLs for a specific connection profile. Remote access VPN administrators can enable or disable the Alias names and Alias URLs. VPN users can choose an Alias name when they connect to the Secure Firewall Threat Defense device. Aliases names for all connections configured on this device can be turned on or off for display. You can also configure the list of Alias URLs, which your endpoints can select while initiating the remote access VPN connection. If users connect using the Alias URL, system will automatically log them using the connection profile that matches the Alias URL.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** on the policy that you want to modify.

Step 3 Click **Edit** on the connection profile for which you want to create or update aliases.

Step 4 Click **Aliases**.

Step 5 To add an Alias name, do the following:

a) Click **Add** under **Alias Names**.
b) Specify the **Alias Name**.
c) Select the **Enabled** check box in each window to enable the aliases.
d) Click **OK**.

Step 6 To add an Alias URL, do the following:

a) Click **Add** under **URL Alias**.
b) Select the **Alias URL** from the list or create a new URL object. For more information see Creating URL Objects, on page 1445.
c) Select the **Enabled** check box in each window to enable the aliases.
d) Click **OK**.

Step 7 Save your changes.

Related Topics

Configure Connection Profile Settings, on page 1594

Configure Access Interfaces for Remote Access VPN

The **Access Interface** table lists the interface groups and security zones that contain the device interfaces. These are configured for remote access SSL or IPsec IKEv2 VPN connections. The table displays the name of each interface group or security-zone, the interface trustpoints used by the interface, and whether Datagram Transport Layer Security (DTLS) is enabled.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select an existing remote access VPN policy in the list and click the corresponding **Edit** icon.

Step 3 Click the **Access Interface** tab.

Step 4 To add an access interface, click + and specify values for the following in the **Add Access Interface** dialog box:

a) **Access Interface**—Select the interface group or security zone to which the interface belongs.

The interface group or security zone must be a Routed type. Other interface types are not supported for remote access VPN connectivity.

b) Associate the **Protocol** object with the access interface by selecting the following options:

- **Enable IPSet-IKEv2**—Select this option to enable **IKEv2** settings.
- **Enable SSL**—Select this option to enable **SSL** settings.
  - Select **Enable Datagram Transport Layer Security**.

    When selected, it enables Datagram Transport Layer Security (DTLS) on the interface and allows the AnyConnect VPN module of Cisco Secure Client to establish an SSL VPN connection using two simultaneous tunnels—an SSL tunnel and a DTLS tunnel.

    Enabling DTLS avoids the latency and bandwidth problems associated with certain SSL connections and improves the performance of real-time applications that are sensitive to packet delays.

To configure SSL settings, and TLS and DTLS versions, see About SSL Settings, on page 980.

To configure SSL settings for the AnyConnect VPN module of Cisco Secure Client, see Group Policy Secure Client Options, on page 1469.

- Select the **Configure Interface Specific Identity Certificate** check box and select **Interface Identity Certificate** from the drop-down list.

  If you do not select the Interface Identity Certificate, the **Trustpoint** will be used by default.

  If you do not select the Interface Identity Certificate or Trustpoint, the **SSL Global Identity Certificate** will be used by default.

c) Click **OK** to save the changes.

Step 5 Select the following under **Access Settings**:

- **Allow Users to select connection profile while logging in**—If you have multiple connection profiles, check this check box to allow user to select the correct connection profile during login. You must select this option for **IPsec-IKEv2** VPNs.
- **Enable HTTP-only VPN Cookies**—Check this check box to enable HTTP-only VPN cookies.

Step 6 Use the following options to configure **SSL Settings**:

- **Web Access Port Number**—The port to use for VPN sessions. The default port is 443.
- **DTLS Port Number**—The UDP port to use for DTLS connections. The default port is 443.
- **SSL Global Identity Certificate**— The selected **SSL Global Identity Certificate** will be used for all the associated interfaces if the **Interface Specific Identity Certificate** is not provided.

Step 7 For **IPsec-IKEv2 Settings**, select the **IKEv2 Identity Certificate** from the list or add an identity certificate.

Step 8 Configure **Service Access Control**. Choose a service access object from the **Service Access Object** drop-down list or click + to create a new object.

You can use a service access object to control remote clients' access to VPN on Firewall Threat Defense devices with Version 7.7 or later. This object provides geolocation-based access control to clients before VPN authentication. By default, there is no access control for RA VPN, and remote clients can connect from any geolocation unless specified by a service access object. For more information, see Manage VPN Access of Remote Users Based on Geolocation, on page 1645and Configure a Service Access Object, on page 1372.

Step 9 Under the **Access Control for VPN Traffic** section, select the following option if you want to bypass access control policy:

- **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** — Decrypted traffic is subjected to Access Control Policy inspection by default. Enabling the Bypass Access Control policy for decrypted traffic option bypasses the ACL inspection, but VPN Filter ACL and authorization ACL downloaded from AAA server are still applied to VPN traffic.

  Note
  If you select this option, you need not update the access control policy for remote access VPN as specified in Update the Access Control Policy on the Secure Firewall Threat Defense Device, on page 1587.

Step 10 Click **Save** to save the access interface changes.

Related Topics

Interface, on page 1398

Configure Advanced Options for Remote Access VPN

Cisco Secure Client Image

Secure Client Image

The Secure Client provides secure SSL or IPsec (IKEv2) connections to the Firewall Threat Defense device for remote users with full VPN profiling to corporate resources. Without a previously-installed client, remote users can enter the IP address of an interface configured to accept clientless VPN connections in their browser to download and install the Secure Client. The Firewall Threat Defense device downloads the client that matches the operating system of the remote computer. After downloading, the client installs and establishes a secure connection. In case of a previously installed client, when the user authenticates, the Firewall Threat Defense device, examines the version of the client, and upgrades the client if necessary.

The Remote Access VPN administrator associates any new or additional Secure Client images to the VPN policy. The administrator can unassociate the unsupported or end of life client packages that are no longer required.

The Secure Firewall Management Center determines the type of operating system by using the file package name. If the user renamed the file without indicating the operating system information, the valid operating system type must be selected from the list box.

Download the Secure Client image file by visiting Cisco Software Download Center.

Related Topics

Adding a Secure Client Image to the Secure Firewall Management Center, on page 1615

Adding a Secure Client Image to the Secure Firewall Management Center

You can upload the Secure Client image to the Secure Firewall Management Center by using the **Secure Client File** object. For more information, see File Objects, on page 1478. For more information about the client image, see Cisco Secure Client Image, on page 1615.

Procedure

Step 1 Choose **Devices** > **Remote Access**, choose and edit a listed remote access policy, then choose the **Advanced** tab.

Step 2 Click **Add** to add a Secure Client image.

Step 3 Click **Add** in the **Available Secure Client Images** portion of the **Secure Client Images** dialog.

Step 4 Enter the **Name** and **Description**(optional) for the available Secure Client Image.

Step 5 Click **Browse**, locate and select the client image that you want to upload.

Step 6 Click **Save** to upload the image to the Firewall Management Center.

When you upload the client image to the Secure Firewall Management Center, the operating system information for the image appears automatically.

Step 7 To change the order of client images, Click **Show Re-order buttons** and move the client image up or down.

Related Topics

Cisco Secure Client Image, on page 1615

Update Secure Client Image for Remote Access VPN Clients

When new Secure Client updates are available in Cisco Software Download Center, you can download the packages manually and add them to the remote access VPN policy so that the new client packages are upgraded on the VPN client systems according to their operating systems.

Before you begin

Instructions in this section help you update new Secure Client images to remote access VPN clients connecting to Secure Firewall Threat Defense VPN gateway. Ensure that the following configurations are complete before updating your Secure Client images:

- Download the latest Secure Client image files from Cisco Software Download Center.
- On your Secure Firewall Management Center web interface, go to **Objects** > **Object Management** > **VPN** > **Secure Client File** and add the new Secure Client image files.

Procedure

Step 1 On your Secure Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** on the remote access VPN policy that you want to update.

Step 3 Click **Advanced** > **Secure Client Image** > **Add**.

Step 4 Select a client image file from **Available Secure Client Images** and click **Add**.

If the required client image is not listed, click **Add** to browse and upload an image.

Step 5 Click **OK**.

Step 6 Save the remote access VPN policy.
After the remote access VPN policy changes are deployed, the new Secure Client images are updated on the Secure Firewall Threat Defense device that is configured as the remote access VPN gateway. When a new VPN user connects to the VPN gateway, the user gets the new Secure Client image to download depending on the operating system of the client system. For existing VPN users, the Secure Client image gets updated in their next VPN session.

Add a Cisco Secure Client External Browser Package to the Secure Firewall Management Center

If you have the Secure Client external browser package image on your local disk, use this procedure to upload the same to the Secure Firewall Management Center. After you upload the external browser package, you can update the external browser package for your remote access VPN connections.

You can upload the external browser package file to the Secure Firewall Management Center by using the **Secure Client File** object. For more information, see File Objects, on page 1478.

Points to Remember

- Only one external browser package can be added to the Firewall Threat Defense device.
- After the external browser package is added to the Firewall Management Center, the browser is pushed to the Firewall Threat Defense only after the external browser is enabled in the remote access VPN configuration.

Procedure

Step 1 On the Secure Firewall Management Center web interface, choose **Devices > Remote Access**, choose and edit a listed remote access policy, then choose the **Advanced** tab.

Step 2 Click **Add** in the **Secure Client External Browser Package** portion of the **Secure Client Images** page.

Step 3 Enter the **Name** and **Description** for the Secure Client package.

Step 4 Click **Browse** and locate the external browser package file to upload.

Step 5 Click **Save** to upload the image to the Secure Firewall Management Center.

Note
If you want to update the remote access VPN connection with an existing external browser package, select the file from the **Package File** drop-down.

Step 6 Save the remote access VPN policy.

Related Topics

Cisco Secure Client Image, on page 1615

Remote Access VPN Address Assignment Policy

The Firewall Threat Defense device can use an IPv4 or IPv6 policy for assigning IP addresses to Remote Access VPN clients. If you configure more than one address assignment method, the Firewall Threat Defense device tries each of the options until it finds an IP address.

IPv4 or IPv6 Policy

You can use the IPv4 or IPv6 policy to address an IP address to remote access VPN clients. You must try with the IPv4 policy to begin and later followed by IPv6 policy.

- **Use Authorization Server**—Retrieves the address from an external authorization server on a per-user basis. If you are using an authorization server that has IP address configured, we recommend using this method. Address assignment is supported by RADIUS-based authorization server only. It is not supported for AD/LDAP. This method is available for both IPv4 and IPv6 assignment policies.
- **Use DHCP**—Obtains IP addresses from a DHCP server configured in a connection profile. You can also define the range of IP addresses that the DHCP server can use by configuring DHCP network scope in the group policy. If you use DHCP, configure the server in the **Objects > Object Management > Network** pane. This method is available for IPv4 assignment policies.

  For more information about DHCP network scope configuration, see Group Policy General Options, on page 1466.

- **Use an internal address pool**—Internally configured address pools are the easiest method of address pool assignment to configure. If you use this method, create the IP address pools in the **Objects > Object Management >Address Pools** pane and select the same in the connection profile. This method is available for both IPv4 and IPv6 assignment policies.

- **Allow reuse an IP address so many minutes after it is released**—Delays the reuse of an IP address after its return to the address pool. Adding a delay helps to prevent problems firewalls can experience when an IP address is reassigned quickly. By default, the delay is set to zero. If you want to extend the delay, enter the number of minutes in the range of 0–480 to delay the IP address reassignment. This configurable element is available for IPv4 assignment policies.

Related Topics

Configure Connection Profile Settings, on page 1594
Remote Access VPN Authentication, on page 1576

Configure Certificate Maps

Certificate maps let you define rules matching a user certificate to a connection profile based on the contents of the certificate fields. Certificate maps provide certificate authentication on secure gateways.

The rules or the certificate maps are defined in Certificate Map Objects, on page 1462.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select an existing remote access VPN policy in the list and click the corresponding **Edit** icon.

Step 3 Choose **Advanced** > **Certificate Maps**.

Step 4 Select the following options from the **General Settings for Connection Profile Mapping** pane:

Selections are priority-based, matching continues down the list of options when the first selection does not match. Matching is complete when the rules are satisfied. If the rules are not satisfied, the default connection profile listed at the bottom of this page is used for the connection. Select any, or all of the following options to establish authentication and to determine which connection profile (tunnel group) must be mapped to the client.

- **Use Group URL if Group URL and Certificate Map match different Connection profiles**

- **Use the configured rules to match a certificate to a Connection Profile**—Enable this to use the rules defined in the Connection Profile Maps.

Note
Configuring a certificate mapping implies certificate-based authentication. The remote user will be prompted for a client certificate regardless of the configured authentication method.

Step 5 Under the **Certificate to Connection Profile Mapping** section, click **Add Mapping** to create certificate to connection profile mapping for this policy.

a) Choose or create a **Certificate Map Name** object.
b) Select the **Connection Profile** that want to use if the rules in the certificate map object are satisfied.
c) Click **OK** to create the mapping.

Step 6 Click **Save**.

Configuring Group Policies

A group policy is a set of attribute and value pairs, stored in a group policy object, that define the remote access VPN experience. For example, in the group policy object, you configure general attributes such as addresses, protocols, and connection settings.

The group policy applied to a user is determined when the VPN tunnel is being established. The RADIUS authorization server assigns the group policy, or it is obtained from the current connection profile.

Note There is no group policy attribute inheritance on the Firewall Threat Defense. A group policy object is used, in its entirety, for a user. The group policy object identified by the AAA server upon login is used, or, if that is not specified, the default group policy configured for the VPN connection is used. The provided default group policy can be set to your default values, but will only be used if it is assigned to a connection profile and no other group policy has been identified for the user.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select an existing remote access VPN policy in the list and click the corresponding **Edit** icon.

Step 3 Choose **Advanced** > **Group Policies** > **Add**.

Step 4 Select group policies from the **Available Group Policy** list and click **Add**. You can select one or more group policies to associate with this remote access VPN policy.

Step 5 Click **OK** to complete the group policy selection.

Step 6 Save your changes.

Related Topics

Configure Group Policy Objects, on page 1466

Configuring LDAP Attribute Mapping

An LDAP attribute name maps LDAP user or group attribute name to a Cisco-understandable name. The attribute map equates attributes that exist in the Active Directory (AD) or LDAP server with Cisco attribute names. You can map any standard LDAP attribute to a well-known vendor specific attribute (VSA). You can map one or more LDAP attributes to one or more Cisco LDAP attributes. When the AD or LDAP server returns authentication to the Firewall Threat Defense device during remote access VPN connection establishment, the Firewall Threat Defense device can use the information to adjust how the Secure Client completes the connection.

When you want to provide VPN users with different access permissions or VPN content, you can configure different VPN policies on the VPN server and assign these policy-sets to each user based on their credentials. You can achieve this in Firewall Threat Defense by configuring LDAP authorization, with LDAP attribute maps. In order to use LDAP to assign a group policy to a user, you must configure a map that maps an LDAP attribute.

An LDAP attribute map consists of three components:

- **Realm**—Specifies the name for the LDAP attribute map; the name is generated based on the selected realm.
- **Attribute Name Map**—Maps the LDAP user or group attribute name to Cisco-understandable name.
- **Attribute Value Map**—Maps value in the LDAP user or group attribute to the value of a Cisco attribute for the selected name mapping.

The group policies that are used in an LDAP attribute map get added to the list of group policies in the remote access VPN configuration. Removing a group policy from the remote access VPN configuration also removes the associated LDAP attribute mapping.

In versions 6.4 to 6.6, you can configure LDAP attribute maps only using FlexConfig. For more information, see Configure AnyConnect Modules and Profiles Using FlexConfig.

In versions 7.0 and later, you can use the following procedure:

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select an existing remote access VPN policy in the list and click the corresponding **Edit** icon.

Step 3 Click **Advanced** > **LDAP Attribute Mapping**.

Step 4 Click **Add**.

Step 5 On the Configure LDAP Attribute Map page, select a **Realm** to configure the attribute map.

Step 6 Click **Add**.

You can configure multiple attribute maps. Each attribute map requires that you configure a name map and value maps.

Note
Ensure that you follow these guidelines while creating an LDAP attribute map:

- Configure at least one mapping for an LDAP attribute; multiple mappings with the same LDAP attribute name is not allowed.
- Configure a minimum of one name map to create an LDAP attribute map.
- You can remove any LDAP attribute map if the attribute map is not associated with any connection profile in the remote access VPN configuration.
- Use the correct spelling and capitalization in the LDAP attribute map for ***both*** the Cisco and LDAP attribute names and values.

a) Specify the **LDAP Attribute Name** and then select the required **Cisco Attribute Name** from the list.
b) Click **Add Value Map** and Specify the **LDAP Attribute Value** and **Cisco Attribute Value**.

Repeat this step to add more value maps.

Step 7 Click **OK** to complete LDAP attribute map configuration.

Step 8 Click **Save** to save the changes to the LDAP attribute mapping.

Example

For a detailed example, see Configure RA VPN with LDAP Authentication and Authorization for FTD.

Related Topics

Configure AAA Settings for Remote Access VPN, on page 1597
Understanding Policy Enforcement of Permissions and Attributes, on page 1578

Configuring VPN Load Balancing

About VPN Load Balancing

VPN load balancing in Firewall Threat Defense allows you group two or more devices logically and distribute remote access VPN sessions among the devices equally. VPN load balancing shares Secure Client VPN sessions among the devices in a load balancing group.

VPN load balancing is based on simple distribution of traffic without taking into account throughput or other factors. A VPN load-balancing group consists of two or more Firewall Threat Defense devices. One device acts as the director, and the other devices are member devices. Devices in a group do not need to be of the exact same type, or have identical software versions or configurations. Any Firewall Threat Defense device that supports remote access VPN can participate in a load balancing group. Firewall Threat Defense supports VPN load balancing with Secure Client SAML authentication.

All active devices in a VPN load-balancing group carry session loads. VPN load balancing directs traffic to the least-loaded device in the group, distributing the load among all devices. It makes efficient use of system resources and provides increased performance and high availability.

Components of VPN Load Balancing

Following are the components of VPN load balancing:

- **Load-balancing group**—A virtual group of two or more Firewall Threat Defense devices to share the VPN sessions.

  A VPN load-balancing group can consist of Firewall Threat Defense devices of the same release or of mixed releases; but the device must support remote access VPN configuration.

  See Configure Group Settings for VPN Load Balancing, on page 1622 and Configure Additional Settings for Load Balancing, on page 1623.

- **Director**—One device from the group acts a director. It distributes the load among other members in the group and participate is serving the VPN sessions.

  The director monitors all devices in the group, keeps track of how loaded each device is, and distributes the session load accordingly. The role of director is not tied to a physical device; it can shift among devices. For example, if the current director fails, one of the member devices in the group takes over that role and immediately becomes the new director.

- **Members**—Devices other than the director in a group are called members. They participate in load balancing and share the remote access VPN connections.

Configure Settings for Participating Devices, on page 1624.

Prerequisites for VPN Load Balancing

- **Certificates**—Firewall Threat Defense's certificate must contain the IP addresses or FQDN of the director and members to which the connection is redirected. Or else, the certificate will be deemed untrusted. The certificate must use Subject Alternate Name (SAN) or wildcard certificate
- **Group URL**—Add the group URL for VPN load-balancing group IP address to the connection profiles. Specify a group URL to eliminate the need for the user to select a group at login.
- **IP Address Pool**—Choose unique IP address pool for member devices, and override the IP pool in Firewall Management Center for each of the member devices.
- Devices that are behind Network Address Translation (NAT) can also be part of a load balancing group.

Guidelines and Limitations for VPN Load Balancing

- VPN load balancing is disabled by default. You must explicitly enable VPN load balancing.
- Only the Firewall Threat Defense devices that are co-located can be added to a load-balancing group.
- A load-balancing group must have a minimum of two Firewall Threat Defense devices.
- Devices in Firewall Threat Defense high availability can participate in a load-balancing group.
- Devices that are behind Network Address Translation (NAT) can also be part of a load balancing group.
- When a member or a director device goes down, remote access VPN connections that are served by that device will be dropped. You must initiate the VPN connection again.
- Identity certificate on each device must have Subject Alternate Name (SAN) or wildcard.

Configure Group Settings for VPN Load Balancing

You can enable VPN load balancing and configure group settings that are applicable to all the members of the load-balancing group. When you create the group, you can configure participation settings for load balancing.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** on the remote access VPN policy that you want to update.

Step 3 Click **Advanced** > **Load Balancing**.

Step 4 Click the **Enable Load balancing between member devices** toggle button to enable load balancing. The **Edit Group Configuration** page opens. Group parameters apply to all devices under the load-balancing group.

Step 5 Specify the **Group IPv4 Address** and **Group IPv6 Address** as applicable.

The IP address that you specify here is for the entire load-balancing group and the director opens this IP address for incoming VPN connections.

Step 6 Select the **Communication Interface** for the load-balancing group. Click **Add** to add an interface group or security zone.

Communication interface is a private interface through which the director and members share information about their load.

Step 7 Enter the **UDP Port** for communication between the director and members in a group. The default port is 9023.

Step 8 Enable the **IPsec Encryption** toggle button to activate IPsec encryption for the communication between the director and members.

Enabling the encryption establishes an IPsec tunnel between the director and members using a pre-shared key.

When you upgrade or downgrade Firewall Threat Defense devices with the **IPsec Encryption** option enabled, ensure there is no configuration mismatch between the Firewall Management Center and the Firewall Threat Defense to prevent deployment failures.

Step 9 Enter **Encryption Key** for IPsec encryption and confirm the encryption key.

Step 10 Click **OK**.

Configure Additional Settings for Load Balancing

The additional settings for VPN load balancing include FQDN and IKEv2 redirection.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** on the remote access VPN policy that you want to update.

Step 3 Click **Advanced** > **Load Balancing**.

Step 4 Turn on the **Enable Load balancing between member devices** toggle button to enable load balancing if not done already.

Step 5 Click **Settings**.

Step 6 Turn on the **Send FQDN to peer devices instead of IP** toggle button to enable redirection using a fully qualified domain name.

By default, Firewall Threat Defense sends only IP addresses in VPN load balancing redirection to a client.

Step 7 Select one of the **IKEv2 Redirect** phases:

- **Redirect during SA authentication**
- **Redirect during SA initialization**

Step 8 Click **OK**.

Step 9 Save your changes.

Configure Settings for Participating Devices

The device participation settings determine how the devices share load in VPN load balancing. Configure a participating device by enabling VPN load balancing on the device and defining device-specific properties. These values vary from device to device. You can provide a priority number for the devices participating in load balancing. A higher priority number gives the device a better chance to become the director over other devices. But you cannot select a device to be the director of the group.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** next to the remote access VPN policy that you want to modify.

Step 3 Click **Advanced** > **Load Balancing**.

Step 4 Turn on the **Enable Load balancing between member devices** toggle button to enable load balancing if you have not enabled already.

Step 5 Configure **Device Participation** settings:

The **Device Participation** section lists all the target devices of the selected remote access VPN configuration. You can configure these devices to share the load of the incoming VPN sessions.

a) Turn on the **Load Balancing** toggle button to enable load balancing for a device and then click **Edit**.
b) Enter the device **Priority**.

By default, the device priority is set to 5. You can choose a number from 1 trough 10.

c) Specify the **IPv4 NAT** or **IPv6 NAT** address for VPN interface IP address if the device is behind NAT.
d) Click **OK**.

Step 6 Click **Save** to save the remote access VPN policy settings.

Configuring IPsec Settings for Remote Access VPNs

The IPsec settings are applicable only if you selected IPsec as the VPN protocol while configuring your remote access VPN policy. If not, you can enable IKEv2 using the Edit Access Interface dialog box. See Configure Access Interfaces for Remote Access VPN, on page 1613 for more information.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 From the list of available VPN policies, select the policy for which you want to modify the settings.

Step 3 Click **Advanced**.

The list of IPsec settings appears in a navigation pane on the left of the screen.

Step 4 Use the navigation pane to edit the following IPsec options:

a) **Crypto Maps**—The Crypto Maps page lists the interface groups on which IKEv2 protocol is enabled. Crypto Maps are auto generated for the interfaces on which IKEv2 protocol is enabled. To edit a Crypto Map, see Configure Remote Access VPN Crypto Maps, on page 1625. You can add or remove interface

groups to the selected VPN policy in **Access Interface**. See Configure Access Interfaces for Remote Access VPN, on page 1613 for more information.

b) **IKE Policy**—The IKE Policy page lists all the IKE policy objects applicable for the selected VPN policy when Secure Client endpoints connect using the IPsec protocol. See IKE Policies in Remote Access VPNs, on page 1627 for more information. To add a new IKE policy, see Configure IKEv2 Policy Objects , on page 1476. Firewall Threat Defense supports only Secure Client IKEv2 clients. Third-party standard IKEv2 clients are not supported.

c) **IPsec/IKEv2 Parameters**—The IPsec/IKEv2 Parameters page enables you to modify the IKEv2 session settings, IKEv2 Security Association settings, IPsec settings, and NAT Transparency settings. See Configure Remote Access VPN IPsec/IKEv2 Parameters, on page 1628 for more information.

Step 5 Click **Save**.

Configure Remote Access VPN Crypto Maps

Crypto maps are automatically generated for the interfaces on which IPsec-IKEv2 protocol has been enabled. You can add or remove interface groups to the selected VPN policy in **Access Interface**. See Configure Access Interfaces for Remote Access VPN, on page 1613 for more information.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 From the list of available VPN policies, select the policy for which you want to modify the settings.

Step 3 Click the **Advanced** > **Crypto Maps**, and select a row in the table and click **Edit** to edit the Crypto map options.

Step 4 Select **IKEv2 IPsec Proposals** and select the transform sets to specify which authentication and encryption algorithms will be used to secure the traffic in the tunnel.

Step 5 Select **Enable Reverse Route Injection** to enable static routes to be automatically inserted into the routing process for those networks and hosts protected by a remote tunnel endpoint.

Step 6 Select **Enable Client Services** and specify the port number.

The Client Services Server provides HTTPS (SSL) access to allow the Secure Client Downloader to receive software upgrades, profiles, localization and customization files, CSD, SCEP, and other file downloads required by the client. If you select this option, specify the client services port number. If you do not enable the Client Services Server, users will not be able to download any of these files that the Secure Client might need.

Note
You can use the same port that you use for SSL VPN running on the same device. Even if you have an SSL VPN configured, you must select this option to enable file downloads over SSL for IPsec-IKEv2 clients.

Step 7 Select **Enable Perfect Forward Secrecy** and select the **Modulus group**.

Use Perfect Forward Secrecy (PFS) to generate and use a unique session key for each encrypted exchange. The unique session key protects the exchange from subsequent decryption, even if the entire exchange was recorded and the attacker has obtained the preshared or private keys used by the endpoint devices. If you select this option, also select the Diffie-Hellman key derivation algorithm to use when generating the PFS session key in the **Modulus Group** list.

Modulus group is the Diffie-Hellman group to use for deriving a shared secret between the two IPsec peers without transmitting it to each other. A larger modulus provides higher security but requires more processing time. The two peers must have a matching modulus group. Select the modulus group that you want to allow in the remote access VPN configuration:

- 1—Diffie-Hellman Group 1 (768-bit modulus).
- 2—Diffie-Hellman Group 2 (1024-bit modulus).
- 5—Diffie-Hellman Group 5 (1536-bit modulus, considered good protection for 128-bit keys, but group 14 is better). If you are using AES encryption, use this group (or higher).
- 14—Diffie-Hellman Group 14 (2048-bit modulus, considered good protection for 128-bit keys).
- 19—Diffie-Hellman Group 19 (256-bit elliptical curve field size).
- 20—Diffie-Hellman Group 20 (384-bit elliptical curve field size).
- 21—Diffie-Hellman Group 21 (521-bit elliptical curve field size).
- 24—Diffie-Hellman Group 24 (2048-bit modulus and 256-bit prime order subgroup).

Step 8 Specify the **Lifetime Duration (seconds)**.

The lifetime of the security association (SA), in seconds. When the lifetime is exceeded, the SA expires and must be renegotiated between the two peers. Generally, the shorter the lifetime (up to a point), the more secure your IKE negotiations will be. However, with longer lifetimes, future IPsec security associations can be set up more quickly than with shorter lifetimes.

You can specify a value from 120 to 2147483647 seconds. The default is 28800 seconds.

Step 9 Specify the **Lifetime Size (kbytes)**.

The volume of traffic (in kilobytes) that can pass between IPsec peers using a given security association before it expires.

You can specify a value from 10 to 2147483647 kbytes. The default is 4,608,000 kilobytes. No specification allows infinite data.

Step 10 Select the following **ESPv3 Settings**:

- **Validate incoming ICMP error messages**—Choose whether to validate ICMP error messages received through an IPsec tunnel and destined for an interior host on the private network.
- **Enable 'Do Not Fragment' Policy**—Define how the IPsec subsystem handles large packets that have the do-not-fragment (DF) bit set in the IP header, and select one of the following from the **Policy** list:
  - Copy—Maintains the DF bit.
  - Clear—Ignores the DF bit.
  - Set—Sets and uses the DF bit.
- Select **Enable Traffic Flow Confidentiality (TFC) Packets**— Enable dummy TFC packets that mask the traffic profile which traverses the tunnel. Use the **Burst**, **Payload Size**, and **Timeout** parameters to generate random length packets at random intervals across the specified SA.

  Note

Enabling traffic flow confidentiality (TFC) packets prevents the VPN tunnel from being idle. Thus the VPN idle timeout configured in the group policy does not work as expected when you enable the TFC packets. See Group Policy Advanced Options, on page 1472.

- Burst—Specify a value from 1 to 16 bytes.
- Payload Size—Specify a value from 64 to 1024 bytes.
- Timeout—Specify a value from 10 to 60 seconds.

Step 11 Click **OK**.

Related Topics

Interface, on page 1398

IKE Policies in Remote Access VPNs

Internet Key Exchange (IKE) is a key management protocol that is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and automatically establish IPsec security associations (SAs). The IKE negotiation comprises two phases. Phase 1 negotiates a security association between two IKE peers, which enables the peers to communicate securely in Phase 2. During Phase 2 negotiation, IKE establishes SAs for other applications, such as IPsec. Both phases use proposals when they negotiate a connection. An IKE proposal is a set of algorithms that two peers use to secure the negotiation between them. IKE negotiation begins by each peer agreeing on a common (shared) IKE policy. This policy states which security parameters are used to protect subsequent IKE negotiations.

Note Firewall Threat Defense supports only IKEv2 for remote access VPNs.

Unlike IKEv1, in an IKEv2 proposal, you can select multiple algorithms and modulus groups in one policy. Since peers choose during the Phase 1 negotiation, this makes it possible to create a single IKE proposal, but consider creating multiple, different proposals to give higher priority to your most desired options. For IKEv2, the policy object does not specify authentication, other policies must define the authentication requirements.

An IKE policy is required when you configure a remote access IPsec VPN.

Configuring Remote Access VPN IKE Policies

The IKE Policy table specifies all the IKE policy objects applicable for the selected VPN configuration when Secure Client endpoints connect using the IPsec protocol. For more information, see IKE Policies in Remote Access VPNs, on page 1627.

Note Firewall Threat Defense supports only IKEv2 for remote access VPNs.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 From the list of available VPN policies, select the policy for which you want to modify the settings.

Step 3 Click **Advanced** > **IKE Policy**.

Step 4 Click **Add** to select from the available IKEv2 policies, or add a new IKEv2 policy and specify the following:

- **Name**—Name of the IKEv2 policy.
- **Description**—Optional description of the IKEv2 policy
- **Priority**—The priority value determines the order of the IKE policy compared by the two negotiating peers when attempting to find a common security association (SA).
- **Lifetime**— Lifetime of the security association (SA), in seconds
- **Integrity**—The Integrity Algorithms portion of the Hash Algorithm used in the IKEv2 policy.
- **Encryption**—The Encryption Algorithm used to establish the Phase 1 SA for protecting Phase 2 negotiations.
- **PRF Hash**—The pseudorandom function (PRF) portion of the Hash Algorithm used in the IKE policy. In IKEv2, you can specify different algorithms for these elements.
- **DH Group**—The Diffie-Hellman group used for encryption.

Step 5 Click **Save**.

Configure Remote Access VPN IPsec/IKEv2 Parameters

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 From the list of available VPN policies, select the policy for which you want to modify the settings.

Step 3 Click **Advanced** > **IPsec**> **IPsec/IKEv2 Parameters**.

Step 4 Select the following for **IKEv2 Session Settings**:

- **Identity Sent to Peers**—Choose the identity that the peers will use to identify themselves during IKE negotiations:
  - **Auto**—Determines the IKE negotiation by connection type: IP address for preshared key, or Cert DN for certificate authentication (not supported).
  - **IP address**—Uses the IP addresses of the hosts exchanging ISAKMP identity information.
  - **Hostname**—Uses the fully qualified domain name (FQDN) of the hosts exchanging ISAKMP identity information. This name comprises the hostname and the domain name.
- **Enable Notification on Tunnel Disconnect**—Allows an administrator to enable or disable the sending of an IKE notification to the peer when an inbound packet that is received on an SA does not match the traffic selectors for that SA. Sending this notification is disabled by default.
- **Do not allow device reboot until all sessions are terminated**—Check to enable waiting for all active sessions to voluntarily terminate before the system reboots. This is disabled by default.

Step 5 Select the following for **IKEv2 Security Association (SA) Settings**:

- **Cookie Challenge**—Whether to send cookie challenges to peer devices in response to SA initiated packets, which can help thwart denial of service (DoS) attacks. The default is to use cookie challenges when 50% of the available SAs are in negotiation. Select one of these options:
  - **Custom**—Specify **Threshold to Challenge Incoming Cookies**, the percentage of the total allowed SAs that are in-negotiation. This triggers cookie challenges for any future SA negotiations. The range is zero to 100%. The default is 50%.
  - Always— Select to send cookie challenges to peer devices always.
  - Never— Select to never send cookie challenges to peer devices.
- **Number of SAs Allowed in Negotiation**—Limits the maximum number of SAs that can be in negotiation at any time. If used with Cookie Challenge, configure the cookie challenge threshold lower than this limit for an effective cross-check. The default is 100 %.
- **Maximum number of SAs Allowed**—Limits the number of allowed IKEv2 connections.

Step 6 Select the following for **IPsec Settings**:

- **Enable Fragmentation Before Encryption**—This option lets traffic travel across NAT devices that do not support IP fragmentation. It does not impede the operation of NAT devices that do support IP fragmentation.
- **Path Maximum Transmission Unit Aging**—Check to enable PMTU (Path Maximum Transmission Unit) Aging, the interval to Reset PMTU of an SA (Security Association).
- **Value Reset Interval**—Enter the number of minutes at which the PMTU value of an SA (Security Association) is reset to its original value. Valid range is 10 to 30 minutes, default is unlimited.

Step 7 Select the following for **NAT Settings**:

- **Keepalive Messages Traversal**—Select whether to enable NAT keepalive message traversal. NAT traversal keepalive is used for the transmission of keepalive messages when there is a device (middle device) located between a VPN-connected hub and spoke, and that device performs NAT on the IPsec flow. If you select this option, configure the interval, in seconds, between the keepalive signals sent between the spoke and the middle device to indicate that the session is active. The value can be from 10 to 3600 seconds. The default is 20 seconds.
- **Interval**—Sets the NAT keepalive interval, from 10 to 3600 seconds. The default is 20 seconds.

Step 8 Click **Save**.

Customize Cisco Secure Client

You can configure the Secure Client customizations using the management center and deploy them to the VPN headend. The threat defense device distributes these customizations to the endpoint when a user connects from the Secure Client. Administrators can customize the following elements at the endpoint:

Table 84: Secure Client Customizations Available in Management Center

Customization	Description	More Info
GUI text and messages	Customize or localize Secure Client GUI text and informational/error messages.	Customize and Localize Secure Client GUI Text and Messages, on page 1632
Icons and images	Customize the logo, images, or icons of the Secure Client GUI.	Customize Secure Client Icons and Images, on page 1633
Scripts	Deploy scripts on the endpoint devices when a client establishes or disconnects a VPN session.	Deploy Scripts on Endpoint Devices Using Secure Client, on page 1634
Binaries	Deploy custom applications using Secure Client APIs.	Deploy Custom Applications Using Cisco Secure Client APIs, on page 1636
Custom Installer Transforms	Customize the Secure Client installer.	Customize the Secure Client Installer, on page 1637
Localized Installer Transforms	Localize the Secure Client installer.	Localize the Client Installer, on page 1637

Guidelines and Limitations for Secure Client Customizations

Guidelines

- Ensure that you import the xml file to the object defined as `DfltCustomization`. When you import the customization object, the Firewall Management Center checks the XML code for validity.

  For example:

  ```
  hostname# import webvpn customization DfltCustomization
  disk0:/csm/defaultcustomizationwebvpn.xml
  ```

- Add the Windows Cisco Secure Client headend deployment package to the remote access VPN policy before deploying customization commands. This package must be included when supporting Windows, MacOS, or Linux clients.

General Limitations

- If there are any customizations earlier than 7.4 configured directly on the threat defense using the CLI, the management center removes these customizations during deployment.
- No clustering support.

Table 85: Limitations for Secure Client Customizations

Customization	Limitations

GUI Text and Message Customizations	• You must restart Secure Client before you use this customization.
	• No support for right-to-left languages.
	• Some strings are truncated in the GUI due to hardcoded field lengths.
	• Some messages are hardcoded in the client. For example:
	  • Status messages (during an update)
	  • Untrusted server messages
	  • Deferred update messages
	• Localization is version-specific.
	If an Administrator created a translation table based on a template of an old Secure Client version, the new messages do not appear to the remote users. The Administrator must merge the latest template with the translation table so that the table has the new messages. You can use third-party tools like Get text to perform the merge.

Icons and Images Customizations	• You must restart Secure Client before you use this customization. • The customization object names must match the Secure Client GUI filenames. File names are different for each operating system and are case-sensitive. For more information about the filenames, extensions, and sizes for each OS, see the Cisco Secure Client Administrator Guide. • Images do not appear properly if the image size is not correct. The management center or the threat defense device does not verify the image size. • MacOS does not support customization of Secure Client images and icons.
Deploy Customized Scripts	• Secure Client runs only one OnConnect and one OnDisconnect script; however, these scripts may launch other scripts. • Scripts can only execute functions that the user has the right to invoke. • You cannot launch the OnConnect script from the Start Before Logon (SBL) GUI.
Deploy Custom Applications Using Cisco Secure Client APIs (Binaries)	• After you use this customization, if you deploy an updated version of the Secure Client on the management center, the client will download the update and replace your custom UI.

Customization	Limitations
Customize the Client Installer	• This customization is available only for Windows.

Customize and Localize Secure Client GUI Text and Messages

You can customize the Secure Client GUI text and the informational/error messages. You also can customize the GUI text and all messages to appear in a preferred language.

Customize GUI Text and Messages

To customize the GUI text or messages, you can edit the messages in the message file. You can update the messages to include more details in the error messages. For example:

- Modify any label in the login dialog box such as **Password** to **Domain Password**.
- Add your support contact details in the error messages.

Localize GUI Text and Messages

Threat defense devices use translation tables to translate the labels and user messages displayed in the Secure Client. When a user connects to the remote access VPN, the Secure Client identifies the locale set on the endpoint and downloads the translation file. Threat defense devices support 128 locales. By default, Secure Client is installed in English.

- For Cisco Secure Client 5.0, the default localization files for various languages are part of the application.
- For AnyConnect Clients 4.x, you can download the localization files for a few languages from Cisco.com and upload them to the management center.

How to Customize Secure Client GUI Text and Messages

Before you begin

Configure one or more remote access VPN policies.

Procedure

Step 1 Get the base template or translation file either from the Secure Client package or from Cisco.com. For example, AnyConnect.po.

Step 2 Use a text editor to add the translations or customize the labels or messages.

Step 3 Update the **msgstr** string corresponding to each **msgid**.

Step 4 Save the file.

Step 5 Create a new Secure Client Customization object.

a) Choose **Objects > Object Management > VPN > Secure Client Customization**.
b) Click **Add Secure Client Customization**.
c) Enter the name and description for the customization.
d) From the **Customization Type** drop-down list, choose **GUI Text and Messages**.
e) From the **Language** drop-down list, choose a language for which you are adding the translation.

f) Click **Browse** and select the translation file. The supported file extensions are .po, .mo, and .txt.

Step 6 Add the customization to the remote access VPN policy.

a) Choose **Devices** > **Remote Access**.
b) Click the edit icon of the remote access VPN policy.
c) Click **Advanced** > **Secure Client Customizations** > **Localization**.
d) Click **+** to select the translation file.
e) Click **Add**.
f) Click **OK**.

Step 7 Click **Save**.

Step 8 Choose **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

What to do next

Verify the Secure Client customizations. For more information, see Verify Secure Client Customizations, on page 1638.

Customize Secure Client Icons and Images

You can customize the logo, images, and icons of the Secure Client GUI using the management center.

To customize the logo, images, and icons of the Secure Client GUI, you must create a Secure Client customization object in the management center. The name of this customization object must match the Secure Client GUI filename. File names are different for each operating system and are case-sensitive. For more information about the filenames, extensions, and sizes for each OS, see the Cisco Secure Client Administrator Guide.

For example, if you want to replace the corporate logo for Windows clients, you must create a customization object with the name as company_logo and add this customization object to the remote access VPN policy. If you use a different name for the customization object, the Secure Client installer does not change the component. However, if you deploy your own executable to customize the GUI, the customization object can have any name.

Location of all images files:

- Windows: **%PROGRAMFILES%\Cisco\Cisco Secure Client\res**
- Linux: **/opt/cisco/secure client/pixmaps**
- macOS: Not supported

How to Customize Secure Client Images and Icons

Before you begin

Configure one or more remote access VPN policies.

Procedure

Step 1 Create the icons or images using the correct extensions, and sizes.

Images do not appear properly if the image size is not correct. The management center or the threat defense device does not verify the image size.

For more information about the filenames, extensions, and sizes, see the Cisco Secure Client Administrator Guide.

Step 2 Create a new Secure Client Customization object.

a) Choose **Objects > Object Management > VPN > Secure Client Customization**.
b) Click **Add Secure Client Customization**.
c) Enter the name and description for the customization object.

Ensure that the customization object name matches the Secure Client GUI filenames. For more information about the filenames, see the Cisco Secure Client Administrator Guide.

d) From the **Customization Type** drop-down list, choose **Icon and Images**.
e) From the **Platform** drop-down list, choose a platform.
f) Click **Browse** and select the file. The supported file extensions are .png, .ico, and .jpeg.
g) Repeat the steps from 2a - f to add multiple icons and images.

Step 3 Add the customization object to the remote access VPN policy.

a) Choose **Devices > Remote Access**.
b) Click the edit icon of the remote access VPN policy.
c) Click **Advanced > Secure Client Customizations > Icons and Images**.
d) Click **+** to select the file.
e) Click **Add**.
f) Click **OK**.

Step 4 Click **Save**.

Step 5 Choose **Deploy > Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

What to do next

Verify the Secure Client customizations. For more information, see Verify Secure Client Customizations, on page 1638.

Deploy Scripts on Endpoint Devices Using Secure Client

Secure Client lets you download and run scripts when the following events occur:

- Establishment of a new client VPN session with the threat defense. The script triggered by this event is an OnConnect script because it requires this filename prefix. Reconnection of the VPN session does not trigger this script.
- Disconnection of a client VPN session with the threat defense. The script triggered by this event is an OnDisconnect script because it requires this filename prefix.

These scripts run asynchronously and do not delay the connection establishment or disconnection. They can be of any extension and must be executable on the endpoint. Secure Client identifies the OnConnect and OnDisconnect scripts by the filename. It looks for a file whose name begins with OnConnect or OnDisconnect regardless of the file extension.

Some examples of this feature are:

- Refresh the group policy upon VPN connection.
- Mount a network drive upon a VPN connection.
- Unmount a network drive upon a VPN disconnection.

To enable scripts, check the **Enable Scripting** option in the VPN profile. By default, the client does not launch scripts. The client does not require the script to be written in a specific language. It requires an application that can run the script to be installed on the client computer. For the client to launch the script, the script must run from the command line.

Secure Client can only launch scripts after the user logs in and establishes a VPN session. You cannot launch the OnConnect script from the Start Before Logon (SBL) GUI. You must check the **Enable Post SBL On Connect Script** option in the VPN profile to trigger the scripts once the user logs in. Secure Client is a 32-bit application. When the scripts run on a 64-bit Windows version, they use the 32-bit version of cmd.exe.

How to Add Customized Scripts for Secure Client

Before you begin

1. Configure a remote access VPN.
2. Enable scripting in the VPN profile.
3. Add the VPN profile to the remote access VPN group policy.

Procedure

Step 1 Create OnConnect and OnDisconnect scripts for a platform.

Step 2 Create a new Secure Client Customization object.

a) Choose **Objects > Object Management > VPN > Secure Client Customization**.
b) Click **Add Secure Client Customization**.
c) Enter the name and description for the customization.
d) From the **Customization Type** drop-down list, choose **Scripts**.
e) From the **Platform** drop-down list, choose a platform.
f) Select one of the following:

- **On Connect**: To select an OnConnect script.
- **On Disconnect**: To select an OnDisconnect script.

g) Click **Browse** and select the script you want to execute on the endpoint.

Step 3 Add the customization to the remote access VPN policy.

a) Choose **Devices > Remote Access**.

b) Click the edit icon of the remote access VPN policy.
c) Click **Advanced > Secure Client Customizations > Localization**.
d) Click **+** to select the translation file.
e) Click **Add**.
f) Click **OK**.

Step 4 Click **Save**.

Step 5 Choose **Deploy > Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

What to do next

Verify the Secure Client customizations. For more information, see Verify Secure Client Customizations, on page 1638.

Deploy Custom Applications Using Cisco Secure Client APIs

For Windows, Linux, or MacOs machines, you can create and deploy a custom client that uses the Secure Client APIs. You can use this client binary file to replace the Secure Client GUI or CLI binary files.

Your executable can call any resource files such as logo images that you import to the management center. When you deploy your own executable, you can use any filenames for your resource files.

The following table lists the filenames of the client executable files for the different operating systems:

Table 86: Filenames of the Secure Client Executable Files

Client OS	Client GUI File	Client CLI File
Windows	vpnui.exe	vpncli.exe
Linux	vpnui	vpn
MacOS	Not supported by the management center deployment. However, you can deploy an executable for the macOS that replaces the client GUI using other means, such as Altiris Agent.	vpn

How to Deploy Custom Applications Using Cisco Secure Client API

Before you begin

Configure one or more remote access VPN policies.

Procedure

Step 1 Create the custom application using the Cisco Secure Client APIs.

Step 2 Create a new Secure Client Customization object.

a) Choose **Objects > Object Management > VPN > Secure Client Customization**.

b) Click **Add Secure Client Customization**.
c) Enter the name and description for the customization.
d) From the **Customization Type** drop-down list, choose **Binary**.
e) From the **Platform** drop-down list, choose a platform.
f) Click **Browse** and select the custom application.

Step 3 Add the customization to the remote access VPN policy.

a) Choose **Devices** > **Remote Access**.
b) Click the edit icon of the remote access VPN policy.
c) Click **Advanced** > **Secure Client Customizations** > **Binaries**.
d) Click + to select the translation file.
e) Click **Add**.
f) Click **OK**.

Step 4 Click **Save**.

Step 5 Choose **Deploy** > **Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Customize the Secure Client Installer

You can customize the Secure Client GUI by creating a custom transform that deploys with the client installer program. You can import the transform to the management center and deploy it to the threat defense device. The threat defense deploys this transform to the endpoint using the client installer program.

Note This customization is available only for Windows.

Localize the Client Installer

You can translate the messages displayed by the Cisco Secure Client installer. The management center uses a transform to translate the messages displayed by the installer. The transform alters the installation but leaves the original security-signed MSI intact. These transforms only translate the installer screens and do not translate the client GUI screens.

Every release of Cisco Secure Client includes a localized transform that administrators can upload to the management center whenever they upload Cisco Secure Client packages with new software. If you use our localization transform, ensure that you update them with the latest release from Cisco.com whenever you upload a new Cisco Secure Client package.

How to Customize or Localize the Client Installer

Before you begin

Configure one or more remote access VPN policies.

Procedure

Step 1 Create a customized or localized transform.

Step 2 Create a new Secure Client Customization object.

a) Choose **Objects > Object Management > VPN > Secure Client Customization**.
b) Click **Add Secure Client Customization**.
c) Enter the name and description for the customization.
d) From the **Customization Type** drop-down list, choose **Customized Installer Transform** or **Localized Installer Transform**.
e) From the **Platform** drop-down list, choose a platform.
f) Click **Browse** and select the transform.

Step 3 Add the customization to the remote access VPN policy.

a) Choose **Devices > Remote Access**.
b) Click the edit icon of the remote access VPN policy.
c) Click **Advanced > Secure Client Customizations > Custom Installer Transforms** or **Localized Installer Transforms**.
d) Click **+** to select the transform.
e) Click **Add**.
f) Click **OK**.

Step 4 Click **Save**.

Step 5 Choose **Deploy > Deployment** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Verify Secure Client Customizations

Validate the Deployment

After the deployment is complete, validate the deployment in the management center. Click the **Transcript Details** ( )) icon to verify the commands generated for the customizations.

Examples

1. The example below shows the transcript details of a GUI text and messages customization: Secure Client en_us localization with customized prompts.

```
import webvpn translation-table AnyConnect language en-us disk0:/AnyConnect_en-us.po
```

2. The example below shows the transcript details of customized icons and images: customization of the Secure Client logo to ABC logo.

```
import webvpn AnyConnect-customization type resource platform win name company_logo.png
 disk0:/company_logo.png
```

3. The example below shows the transcript details of the customized OnConnect and OnDisconnect scripts. The On connect script mounts the network drive and the On disconnect script unmounts the network drive.

```
import webvpn AnyConnect-customization type binary platform win name
  scripts_OnConnect_mount.bat disk0:/mount.bat
import webvpn AnyConnect-customization type binary platform win name
  scripts_OnDisconnect_unmount.bat disk0:/unmount.bat
```

Verify Customizations Using Threat Defense Commands

Use the following commands on the threat defense to verify the customizations:

- **show import webvpn translation-table detailed**: Shows the available translation tables.

```
HQ-FTD# show import webvpn translation-table detailed
 Translation Tables' Templates:
   AnyConnect           ia4DaAXNSvl5pZboQRGJcs9KMXY=
   customization
 Translation Tables:
   fr                   customization        BWWodsOt1PbvDvYOp8hLb3W7a64=
   ja                   customization        lNvUk1+qTLNZyNrBcApMQPHnm1M=
   ru                   customization        UqyKyUAcjR+xTGUtdiIFnoIiW5U=
```

- **show import webvpn AnyConnect-customization detailed**: Shows the details of the Secure Client customizations.

```
HQ-FTD# show import webvpn AnyConnect-customization detailed
OEM resources for AnyConnect client:
  linux-64/binary/scripts_OnConnect_conn.sh           w6+n7z80D/8AR+ul2f7DvTmcDTw=
  linux-64/binary/scripts_OnDisconnect_discon.sh      jx5LJC2XBEmEkGeww59CAkszvnI=
  linux-64/resource/company-logo.png                  GsfBDroqGSQEEwuBDS/3DJNVv88=
  win/binary/scripts_OnConnect_mount.bat              dzjfsLYYft/XMlPlzskKl+Wvlbw=
  win/binary/scripts_OnDisconnect_unmount.bat         k6xlKhF1l2IRyJu08+sdYXgKNgM=
  win/resource/company_logo.png                       cmEvxwqvtaS+Pz/6sb9n3NZudS4=
```

Verify Customizations in Secure Client

- In Secure Client, click the **Message History** tab to verify that the customizations are downloaded.

  Use the DART tool to view the client-side diagnostics.

Table 87: Verifying the Customizations in Secure Client

Customization	Verification

GUI text and message customizations	• Verify if Secure Client has the language localization or customized files in the following locations:
	  • Windows: %ProgramData%\Cisco\Cisco AnyConnect Secure Mobility Client\l10n\<language-code>\LC_MESSAGES (AnyConnect versions 4.9 and earlier)
	  • Windows: %ProgramData%\Cisco\Cisco Secure Client\l10n\<language-code>\LC_MESSAGES (Secure Client versions 5.0 and later)
	  • For Mac OS and Linux: /opt/cisco/anyconnect/l10n/<LANGUAGE-CODE>/LC_MESSAGES
	• Verify the contents of the localized or customized file and confirm if they have the customized or localized strings. For example: AnyConnect.mo

Customization	Verification

Image and icon customizations	• Verify if Secure Client has the customized files in the following locations:
	  • Windows: %PROGRAMFILES%\ Cisco\Cisco AnyConnect Secure Mobility Client\res\ (AnyConnect versions 4.10 and earlier)
	  • Windows: %PROGRAMFILES%\ Cisco\Cisco Secure Client\UI\res (Cisco Secure Client 5.0 and above)
	  • Mac OS and Linux: /opt/cisco/anyconnect/res
	• Verify the content of the customized icon or image file.

Customized OnConnect and OnDisconnect scripts	• Verify if the customized scripts are in the following locations:
	  • Windows: %ProgramData%\Cisco\Cisco Secure Client\Script (Cisco Secure Client 5.0 and above)
	  • Windows: %ProgramData%\Cisco\ Cisco AnyConnect Secure Mobility Client\Script (AnyConnect versions 4.9 and earlier)
	  • Mac OS and Linux: /opt/cisco/anyconnect/Script
	• Verify the script.

Configure Secure Client Management VPN Tunnel

A management VPN tunnel provides connectivity to the corporate network whenever a client system is powered up, without the VPN users having to connect to the VPN. This helps organizations keep their endpoints up-to-date with software patches and updates. Management tunnel disconnects when the user-initiated VPN tunnel is established.

This section provides information about configuring Secure Client management VPN tunnel on Firewall Threat Defense. Configuring the Secure Client management tunnel on Firewall Threat Defense using the Firewall Management Center web interface requires the following settings:

- A **Connection profile** with certificate-based authentication and a group URL.
- **Secure Client management VPN profile file**, configured a server with group URL and backup servers if required.
- A **Group policy** with the management VPN profile, split tunneling with explicitly included networks, client bypass protocol, and no banner.

For detailed instructions on configuring the Secure Client Management VPN tunnel, see Configuring Secure Client Management VPN Tunnel on Firewall Threat Defense, on page 1641.

Requirements and Prerequisites for Secure Client Management VPN Tunnel

Software and Configuration Requirements

Ensure that you have the following before you configure the Secure Client Management tunnel on using the Firewall Threat Defense using the Firewall Management Center web interface:

- Ensure that you are using Firewall Threat Defense and Firewall Management Center versions 6.7.0 or above.
- Download the Secure ClientSecure Client VPN Webdeploy package 4.7 or above and upload it to Firewall Threat Defense remote access VPN.
- Ensure that the certificate authentication is configured in the connection profile.
- Ensure that no banner is configured in the group policy.
- Check the split tunneling configuration in the management tunnel-group policy.

Certificate Requirements

- Firewall Threat Defense must have a valid identity certificate for remote access VPN and the root certificate from the local certifying authority (CA) must be present on the Firewall Threat Defense.
- Endpoints connecting to the management VPN tunnel must have a valid identity certificate.
- CA certificate for Firewall Threat Defense's identity certificate must be installed on the endpoints and the CA certificate for the endpoints must be installed on the Firewall Threat Defense.
- The identity certificate issued by the same local CA must be present in the Machine store.

  Certificate Store (For Windows) and/or in System Keychain (For macOS).

Limitations of Secure Client Management VPN Tunnel

- Secure Client Management VPN Tunnel supports only certificate authentication, it does not support AAA-based authentication.
- Public or private proxy settings are not supported.
- Secure Client upgrade and AnyConnect module download are not supported when the management VPN tunnel is connected.

Configuring Secure Client Management VPN Tunnel on Firewall Threat Defense

Procedure

Step 1 **Create a remote access VPN policy configuration using the wizard**:

For information about configuring a remote access VPN, see Configuring a New Remote Access VPN Connection, on page 1583.

Step 2 **Configure connection profile settings for management VPN tunnel**:

Note
It is advisable to create a new connection profile to be used only for Secure Client management VPN tunnel.

a) Edit the remote access VPN policy you have created.
b) Select and edit the connection profile that will be used for management VPN tunnel.
c) Click **AAA** > **Authentication Method** and select **Client Certificate Only**. Configure the authorization and accounting settings as required.
d) Click the **Aliases** tab of the connection profile.
e) Click **Add (+)** under URL Aliases and **URL Alias** for the connection profile.
f) Click **Enabled** to enable the URL.
g) Click **OK** and then click **Save** to save the connection profile settings.

For more information about connection profile settings, see Configure Connection Profile Settings, on page 1594.

Step 3 **Create a management tunnel profile using the Secure Client profile editor**:

a) Download the Secure Client **VPN Management Tunnel Standalone Profile Editor** from Cisco Software Download Center if you have not done already.
b) Create a management tunnel profile with the required settings for your VPN users and save the file.
c) Configure a server in the Server List with the group URL you have configured in the connection profile.

For information about creating a management profile using the Profile Editor, see the *Cisco Secure Client (including AnyConnect) Administrator Guide*.

Step 4 **Create a management tunnel object**:

a) On your Secure Firewall Management Center web interface, navigate to **Object** > **Object Management** > **VPN** > **Secure Client File**
b) Click **Add Secure Client File**.
c) Specify the **Name** for the Secure Client file.
d) Click **Browse** and select the management tunnel profile file you have saved.
e) Click the **File Type** drop-down and select **Secure Client Management VPN Profile**.
f) Click **Save**.

Note
You an also create the management tunnel object when you create or update Secure Client settings for a group policy. See Group Policy Secure Client Options, on page 1469.

Step 5 **Associate a management profile with a group policy and configure group policy settings**:

You must add the management VPN profile to the group policy associated with the connection profile used for the management tunnel VPN connection. When the user connects, the management VPN profile is downloaded along with the user VPN profile already mapped to the group policy, enabling the management VPN tunnel feature.

Caution
No Banner: Check and ensure that no banner is configured in the group policy settings. You can check the banner settings under **Group Policy** > **General Settings** > **Banner**.

a) Edit the connect profile you have created for management VPN tunnel.
b) Click **Edit Group Policy** > **Secure Client** > **Management Profile**.
c) Click the **Management VPN Profile** drop-down and select the management profile file object you have created.

Note
You can also click + and add a new Secure Client Management VPN Profile object.

d) Click **Save**.

Step 6 **Configure split tunneling in group policy**:

a) Click **Edit Group Policy** > **General** > **Split Tunneling**.
b) From the IPv4 or IPv6 split tunneling drop-down, select **Tunnel networks specified below**.
c) Select the Split Tunnel Network List Type: **Standard Access List** or **Extended Access List**, and then select the required access list to allow the traffic over the management VPN tunnel.
d) Click **Save** to save the split tunnel settings.

Secure Client Custom Attribute

Secure Client Management VPN tunnel requires split include tunneling configuration by default. If you are configuring Secure Client custom attribute in the group policy to deploy the management VPN tunnel with split tunneling to tunnel all, you can do so using FlexConfig because Firewall Management Center 6.7 web interface does not support Secure Client custom attribute.

The following is an example command for Secure Client custom attribute:

```
webvpn
 anyconnect-custom-attr ManagementTunnelAllAllowed description ManagementTunnelAllAllowed
anyconnect-custom-data ManagementTunnelAllAllowed true true
group-policy MGMT_Tunnel attributes
 anyconnect-custom ManagementTunnelAllAllowed value true
```

Step 7 **Deploy, verify, and monitor the remote access VPN policy**:

a) Deploy the management VPN tunnel configuration to Firewall Threat Defense.

Note
Client systems must connect to the Firewall Threat Defense remote access VPN once to download the management tunnel VPN profile to the client machines.

b) You can verify the Secure Client management VPN tunnel at **Secure Mobility Client** > **VPN** > **Statistics**.

You can also check the management VPN session details on the Firewall Threat Defense command prompt using the **show vpn-sessiondb anyconnect** command.

c) On your Firewall Management Center web interface, click **Analysis** to view the management tunnel session information.

Related Topics

Configure Connection Profile Settings, on page 1594
Firewall Threat Defense Group Policy Objects, on page 1465

Multiple Certificate Authentication

Multiple certificate based authentication gives the ability to have the Firewall Threat Defense validate the machine or device certificate, to ensure the device is a corporate-issued device, in addition to authenticating the user's identity certificate to allow VPN access using the Secure Client during SSL or IKEv2 EAP phase.

The multiple certificates option allows certificate authentication of both the machine and user via certificates. Without this option, you could only do certificate authentication of either machine or the user, but not both.

Guidelines and Limitations of Multiple Certificate Authentication

Note When you configure multiple certificate authentication, ensure that you set the value of **AutomaticCertSelection** to true in the Cisco Secure Client Profile settings.

- Multiple certificate authentication currently limits the number of certificates to two.
- Secure Client must indicate support for multiple certificate authentication. If that is not the case then the gateway uses one of the legacy authentication methods or fails the connection. Secure Client version 4.4.04030 or later supports Multi-Certificate based authentication.
- Secure Client supports only RSA-based certificates.
- Only SHA256, SHA384, and SHA512 based certificate are supported during the Secure Client aggregate authentication.
- Certificate authentication cannot be combined with SAML authentication.

Configuring Multiple Certificate Authentication

Before you begin

Before you configure multiple certificate authentication, ensure that you have configured the certificate enrollment object that is used to obtain the identity certificate for each Firewall Threat Defense device. For more information, see Certificate Map Objects, on page 1462.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select the remote access VPN policy and click **Edit**.

Note
If you have not configured a remote access VPN, click **Add** to create a new remote access VPN policy.

Step 3 Select and **Edit** a connection profile to configure multiple certificate authentication.

Step 4 Click **AAA** settings and select **Authentication Method** > **Client Certificate Only** or **Client Certificate & AAA**.

Note
Select the **Authentication Server** if you have selected the Client Certificate & AAA authentication method

Step 5 Select the **Enable multiple certificate authentication** checkbox.

Step 6 Choose one of the certificates to **Map username from client certificate**:

- **First Certificate**— Select this option to map the username from the machine certificate sent from the VPN client.

- **Second Certificate**— Select this option to map the username from the user certificate sent from the client.

The username sent from the client is used as the VPN session username when certificate only authentication is enabled. When AAA and certificate authentication is enabled, VPN session username will be based on prefill option.

Note
If you select the **Map specific field** option, which includes the username from the client certificate, the **Primary** and **Secondary** fields display default values: **CN (Common Name)** and **OU (Organisational Unit)** respectively.

If you select the **Use entire DN (Distinguished Name) as username** option, the system automatically retrieves the user identity. A distinguished name (DN) is a unique identification, made up of individual fields that can be used as the identifier when matching users to a connection profile DN rules are used for enhanced certificate authentication.

If you have selected the Client Certificate & AAA authentication, select the **Prefill username from certificate on user login window** option to prefill the secondary username from the client certificate when the user connects via AnyConnect VPN module of Cisco Secure Client.

- **Hide username in login window**: The secondary username is pre-filled from the client certificate, but hidden to the user so that the user does not modify the pre-filled username.

Step 7 Configure the required AAA settings and connection profile settings for the remote access VPN.

Step 8 Save the connection profile and remote access VPN configuration and deploy it on your Firewall Threat Defense device.

Related Topics

Configure AAA Settings for Remote Access VPN, on page 1597

Manage VPN Access of Remote Users Based on Geolocation

You can manage the remote access VPN connections of users based on their geolocations. By configuring rules such that they allow or deny VPN connections from specific countries or regions, you can meet compliance requirements and enhance security. Connections that don't meet these location-based criteria are blocked before authentication and the details are logged in the Firewall Management Center (**Devices** > **Troubleshoot** > **Troubleshooting Logs**).

This feature supports all versions of Secure Client.

Prerequisites

Firewall Threat Defense device must be Version 7.7.0 or later.

Workflow to Manage VPN Access of Remote Users Based on Geolocation

Step	Task	More Information
1	Define a policy to allow geolocation-based access control for your remote clients.	Configure a Service Access Object, on page 1372

Step	Task	More Information
2	Update the remote access VPN configuration with the policy.	Configure Access Interfaces for Remote Access VPN, on page 1613
3	Deploy the configurations on the devices.	—
4	After clients connect to the Firewall Threat Defense device: 1. Verify the active remote access VPN sessions in the **Remote Access VPN** dashboard. 2. Verify the logs for the denied remote access VPN sessions at **Devices** > **Troubleshoot** > **Troubleshooting Logs**.	Monitor and Troubleshoot Service Access Policies, on page 1646

Guidelines and Limitations for Managing Remote Access VPN Users Based on Geolocation

Guidelines

- In a service access object, if you use a geolocation object (country, continent, or geolocation object), use it only in one rule.
- Configure the service access rules in the correct order because you cannot reorder these rules.

Limitations

- Clustering is not supported.
- Geolocation-based unclassified IP addresses are not categorized according to their geographic origin. For such unclassified IP addresses, the Firewall Management Center enforces the default service access policy action.
- Connections from IETF RFC1918 addresses are allowed for geolocation-based remote access VPN regardless of service-access policy settings.

Monitor and Troubleshoot Service Access Policies

Monitor Active Remote Access VPN Sessions in Remote Access VPN Dashboard

Choose **Overview** > **Dashboards** > **Remote Access VPN**.

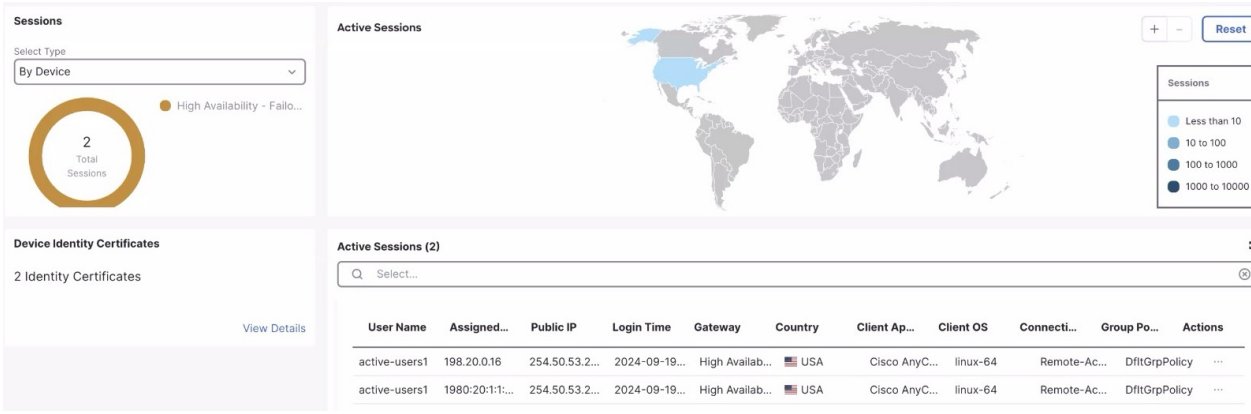

Monitor Denied Remote Access VPN Sessions

Monitor the denied remote access VPN sessions at **Devices** > **Troubleshoot** > **Troubleshooting Logs**. To view the denied remote access VPN sessions, you must configure the syslog settings in the Firewall Threat Defense:

1. Choose **Devices** > **Platform Settings** and create or edit a threat defense policy.
2. In the left pane, click **Syslog**.
3. Click the **Logging Setup** tab.
4. Check the **Enable Logging** check box.
5. Click the **VPN Logs** radio button.
6. From the **Logging Level** drop-down list, choose **6 - informational**.
7. Click **Save**.

Note You cannot view the denied remote access VPN sessions if the **All Logs** option is configured with a **Logging Level** between 0 and 2.

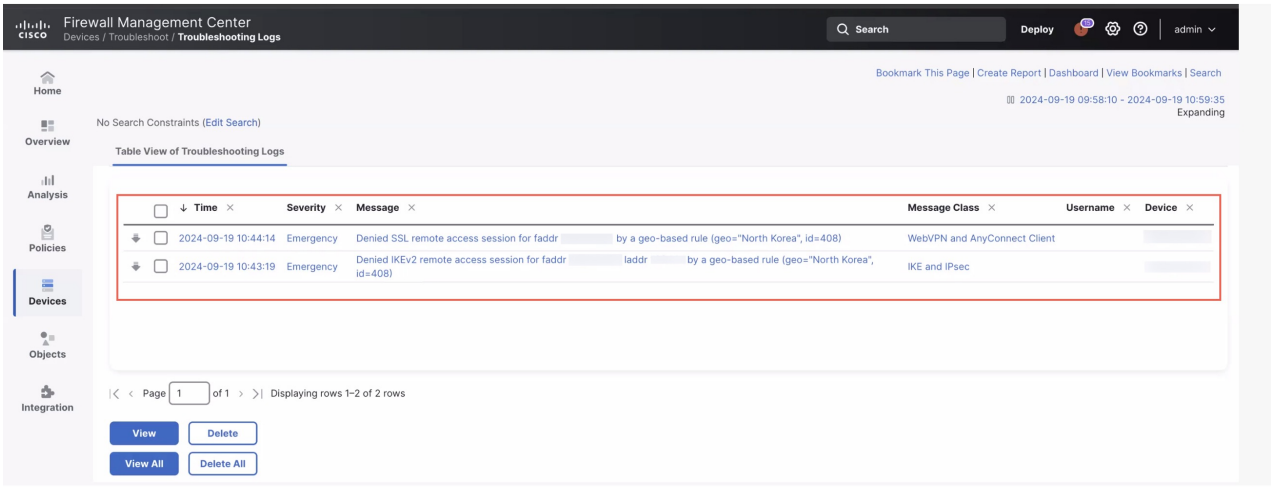

Verify Service Access Policies

From the Firewall Threat Defense device CLI, run the following commands:

- **show running-config service-access**: Displays the user-defined service access policies.

```
firepower#show running-config service-access
   service-access deny geolocation OBJGRP_Asia1
   service-access permit interface outside ra-ssl-client geolocation OBJGRP_India
   service-access deny ra-ikev2 geolocation any
```

- **show service-access**: Displays details of the user-defined service access policies.

```
firepower# show service-access
1 outside          : ra-ikev2 ra-ssl-client (permit) hits = 8288
  Last hit time    : 10:58:10.038 IST Tue Jul 16 2024
  object-group     : FMC_INTERNAL_XXY
2 any              : ra-ikev2 ra-ssl-client (deny) hits = 123
  Last hit time    : 11:23:12.032 IST Tue Jul 17 2024
  object-group     : any


firepower# show service-access detail
1 outside          : ra-ikev2 ra-ssl-client (permit) hits = 8288
  Last hit time    : 10:58:10.038 IST Tue Jul 16 2024
  object-group     : FMC_INTERNAL_XXY
  geolocation      : Egypt(818) Jordan(400)
                     Iran (Islamic Republic of)(364)
                     Saudi Arabia(682)
```

- **show geodb**: Displays details of the geolocation table.

  show geodb { **ipv4** | **ipv6** | **counters** | **context** } [**location** *country_name* | **lookup** *ip_address*] [**detail**]

  - **show geodb** { **ipv4** | **ipv6** } : Displays the total number of IPv4 or IPv6 address mappings.

```
firepower# show geodb ipv4
 Geolocation Table - IPv4
 Total number of mappings available: 532507
 Last geolocation data read time:  17:02:13.000 IST Thu Jul 18 2024
```

```
Running geolocation update version: 2024-02-15-019
```

- **show geodb** { **ipv4** | **ipv6** } **location** *country_name* **detail**: Displays the details of the IPv4 or IPv6 address mappings.

```
firepower# show geodb ipv4 location Antarctica detail
Geolocation Table - IPv4
id=0x00007fff82c284e0, geo_id=10, hits=0
        range_lower=77.70.176.176, range_upper=77.70.176.183
id=0x00007fff82cca360, geo_id:10, hits=0
     range_lower=79.110.169.69, range_upper=79.110.169.69
Total number of mappings available: 28
```

- **show geodb counters**: Displays the details of active, permitted, and denied sessions.

```
firepower# show geodb counters
current   - ongoing sessions
permitted - cumulative permitted sessions
denied    - cumulative denied sessions
Location                current          permitted      denied
Egypt                     0                  0             5
India                     45                 1345         45
```

- **show geodb** { **ipv4** | **ipv6lookup** *ip_address*: Displays the geolocation of a specific IPv4 or IPv6 address.

```
firepower# show geodb ipv4 lookup 223.223.128.24
Geolocation of 223.223.128.24 is "India" (356) with id=0x000015114d0aa330
Matching network range: 223.223.128.0 - 223.223.159.255
```

Troubleshoot Service Access Policies

- **Syslogs**

  Enable remote access VPN service access syslogs:

  1. Choose **Devices** > **Platform Settings**.
  2. Create or edit a platform settings policy.
  3. In the left pane, click **Syslog**.
  4. Click the **Logging Setup** tab and check the **Enable Logging** check box.
  5. Click the **Syslog Settings** tab and enable the syslogs for service access syslog 751031 and 716166.

- **Commands**
  - Use the **show running-config service-access**, and **show service-access** commands to view details of user-defined service access policies.
  - Use the **show geodb** command to view details of the geolocation table.
  - Use the **debug geolocation <debug-level>** command to capture debug logs related to geolocation. The debug levels can be 1 (Error), 2 (Warning), 3 and 4 (Info), 5 (Debug), or 255 (Debug all).

- Use the **clear geodb counters** command to clear the geolocation table counters such as the hit counts of the service access policies. However, you cannot clear the actual permitted and denied counters for locations using this command; you can clear these counters only after you reboot the device.

Customizing Remote Access VPN AAA Settings

This section provides information about customizing your AAA preferences for remote access VPNs. For more information, see Configure AAA Settings for Remote Access VPN, on page 1597.

Authenticate VPN Users via Client Certificates

You can configure remote access VPN authentication using client certificate when you create a new remote access VPN policy using the wizard or by editing the policy later.

Before you begin

Configure the certificate enrollment object that is used to obtain the identity certificate for each Firewall Threat Defense device that acts as a VPN gateway.

Procedure

Step 1 On your Secure Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select a remote access policy and click **Edit**; or click **Add** to create a new remote access VPN policy.

Step 3 For a new remote access VPN policy, configure the authentication while selecting connection profile settings. For an existing configuration, select the connection profile that includes the client profile, and click **Edit**.

Step 4 Click **AAA** > **Authentication Method** > **Client Certificate Only**.

With this authentication method, the user is authenticated using a client certificate. You must configure the client certificate on VPN client endpoints. By default, the user name is derived from client certificate fields CN and OU respectively. If the user name is specified in other fields in the client certificate, use 'Primary' and 'Secondary' field to map appropriate fields.

If you select **Map specific field** option, which includes the username from the client certificate. The **Primary** and **Secondary** fields display the following default values, respectively: **CN (Common Name)** and **OU (Organisational Unit)**. If you select the **Use entire DN as username** option, the system automatically retrieves the user identity. A distinguished name (DN) is a unique identification, made up of individual fields, that can be used as the identifier when matching users to a connection profile. DN rules are used for enhanced certificate authentication.

- Primary and Secondary fields pertaining to the **Map specific field** option contain these common values:
  - C (Country)
  - CN (Common Name)
  - DNQ (DN Qualifier
  - EA (Email Address)

- GENQ (Generational Qualifier)
- GN (Given Name)
- I (Initial)
- L (Locality)
- N (Name)
- O (Organisation)
- OU (Organisational Unit)
- SER (Serial Number)
- SN (Surname)
- SP (State Province)
- T (Title)
- UID (User ID)
- UPN (User Principal Name)

- Whichever authentication method you choose, select or deselect **Allow connection only if user exists in authorization database**.

For more information, see Configure AAA Settings for Remote Access VPN, on page 1597.

Step 5 Save your changes.

Related Topics

Configure Connection Profile Settings, on page 1594
Adding Certificate Enrollment Objects, on page 1414

Configure VPN User Authentication via Client Certificate and AAA Server

When you configure remote access VPN authentication to use both client certificate and authentication server, VPN client authentication is done using both the client certificate validation and AAA server.

Before you begin

- Configure the certificate enrollment object that you use to obtain the identity certificate for each Firewall Threat Defense device that acts as a VPN gateway.
- Configure the RADIUS server group object and any AD or LDAP realms to use in the remote access VPN policy configuration.
- Ensure that the AAA Server is reachable from the Secure Firewall Threat Defense device for the remote access VPN configuration to work.

Procedure

Step 1 On your Secure Firewall Management Center web interface, choose **Devices** > **Remote Access**.

Step 2 Click **Edit** on the remote access VPN policy for which you want to update the authentication or click **Add** to create new one.

Step 3 If you choose to create new remote access VPN policy, configure the authentication while selecting connection profile settings. For an existing configuration, select the connection profile that includes the client profile, and click **Edit**.

Step 4 Go to **AAA** and from the **Authentication Method** drop-down, choose **Client Certificate & AAA**.

- When you select the **Authentication Method** as:

  Client Certificate & AAA—Both types of authentication are done.

  - **AAA**—If you select the **Authentication Server** as **RADIUS**, by default, the Authorization Server has the same value. Select the **Accounting Server** from the drop-down list. Whenever you select **AD** and **LDAP** from the Authentication Server drop-down list, you must manually select the **Authorization Server** and **Accounting Server** respectively.
  - **Client Certificate**—Authenticates the user with client certificate. You must configure client certificate on the VPN client endpoints. By default, the username is derived from client certificate fields CN & OU respectively. If you use any other field in the client profile to specify the username, use **Primary Field** and **Secondary Field** to map appropriate fields.

    If you select **Map specific field** option, which includes the username from the client certificate. The **Primary** and **Secondary** fields display default values: **CN (Common Name)** and **OU (Organisational Unit)** respectively. If you select the **Use entire DN as username** option, the system automatically retrieves the user identity. A distinguished name (DN) is a unique identification, made up of individual fields that can be used as the identifier when matching users to a connection profile. DN rules are used for enhanced certificate authentication.

    Primary and Secondary fields pertaining to the **Map specific field** option contains these common values:

    - C (Country)
    - CN (Common Name)
    - DNQ (DN Qualifier
    - EA (Email Address)
    - GENQ (Generational Qualifier)
    - GN (Given Name)
    - I (Initial)
    - L (Locality)
    - N (Name)
    - O (Organisation)
    - OU (Organisational Unit)

- SER (Serial Number)
- SN (Surname)
- SP (State Province)
- T (Title)
- UID (User ID)
- UPN (User Principal Name)

- Whichever authentication method you choose, select or deselect **Allow connection only if user exists in authorization database**.

For more information, see Configure AAA Settings for Remote Access VPN, on page 1597.

Step 5 Save your changes.

Related Topics

Configure Connection Profile Settings, on page 1594
Adding Certificate Enrollment Objects, on page 1414

Manage Password Changes over VPN Sessions

Password management allows remote access VPN policy administrator to configure the notification settings for the remote access VPN users on their password expiry. Password management is available in AAA settings with authentication methods AAA Only and Client Certificate & AAA. For more information, see Configure AAA Settings for Remote Access VPN, on page 1597.

Procedure

Step 1 On your Secure Firewall Management Center web interface, choose **Devices** > **Remote Access**.

Step 2 Click **Edit** on the remote access VPN policy that you want to update.

Step 3 Click **Edit** on the connection profile that includes AAA settings.

Step 4 Choose **AAA** > **Advanced Settings** >.

Step 5 Check the **Enable Password Management** check-box and select one of the following:

- Notify User - days ahead of password expiry and specify the number of days in the box.
- Notify user on the day of password expiration.

Step 6 Save your changes.

Related Topics

Configure Connection Profile Settings, on page 1594

Send Accounting Records to the RADIUS Server

Accounting records in remote access VPN help the VPN administrator track the services that users access and the amount of network resources that they consume. Accounting information includes when user session start and stop, username, the number of bytes that pass through the device for each session, the service used, and the duration of each session.

You can use accounting alone or together with authentication and authorization. When you activate AAA accounting, the network access server reports the user activity to the configured accounting server. You can configure a RADIUS server as the accounting server so that the Firewall Management Center sends all the user activity information to the RADIUS server.

Note You can use the same RADIUS server or separate RADIUS servers for authentication, authorization, and accounting in remote access VPN AAA settings.

Before you begin

- Configure a RADIUS group object with RADIUS servers to receive authentication requests or accounting records. For more information, see RADIUS Server Group Options, on page 1364.
- Ensure that the RADIUS server is reachable from the Firewall Threat Defense device. Configure routing on your Secure Firewall Management Center at **Devices** > **Device Management** > **Edit Device** > **Routing** to ensure connectivity to the RADIUS server.

Procedure

Step 1 On your Secure Firewall Management Center web interface, choose **Devices** > **Remote Access**.

Step 2 Click **Edit** on the remote access policy for which you want to configure RADIUS server, or create new remote access VPN policy.

Step 3 Click **Edit** on the connection profile that includes AAA settings and choose **AAA**.

Step 4 Select the RADIUS server from the **Accounting Server** drop-down.

Step 5 Save your changes.

Related Topics

Configure Connection Profile Settings, on page 1594
Configure AAA Settings for Remote Access VPN, on page 1597

Delegating Group Policy Selection to Authorization Server

The group policy applied to a user is determined when the VPN tunnel is being established. You can select a group policy for a connection profile while creating a remote access VPN policy using the wizard or update the connection policy for connection profiles later. However, you can configure the AAA (RADIUS) server to assign the group policy or it is obtained from the current connection profile. If the Firewall Threat Defense device receives attributes from the external AAA server that conflicts with those configured on the connection profile, then attributes from the AAA server always take the precedence.

You can configure ISE or the RADIUS Server to set the Authorization Profile for a user or user-group by sending IETF RADIUS Attribute 25 and map to the corresponding group policy name. You can configure specific group policy to a user or user group to push a Downloadable ACL, set a banner, Restrict VLAN, and configure the advanced option of applying an SGT to the session. These attributes are applied to all users that are part of that group when the VPN connection is established.

For more information, see the Configure Standard Authorization Policies section of Cisco Identity Services Engine Administrator Guide and RADIUS Server Attributes for Secure Firewall Threat Defense, on page 1602.

Figure 435: Remote Access VPN Group Policy Selection by AAA Server

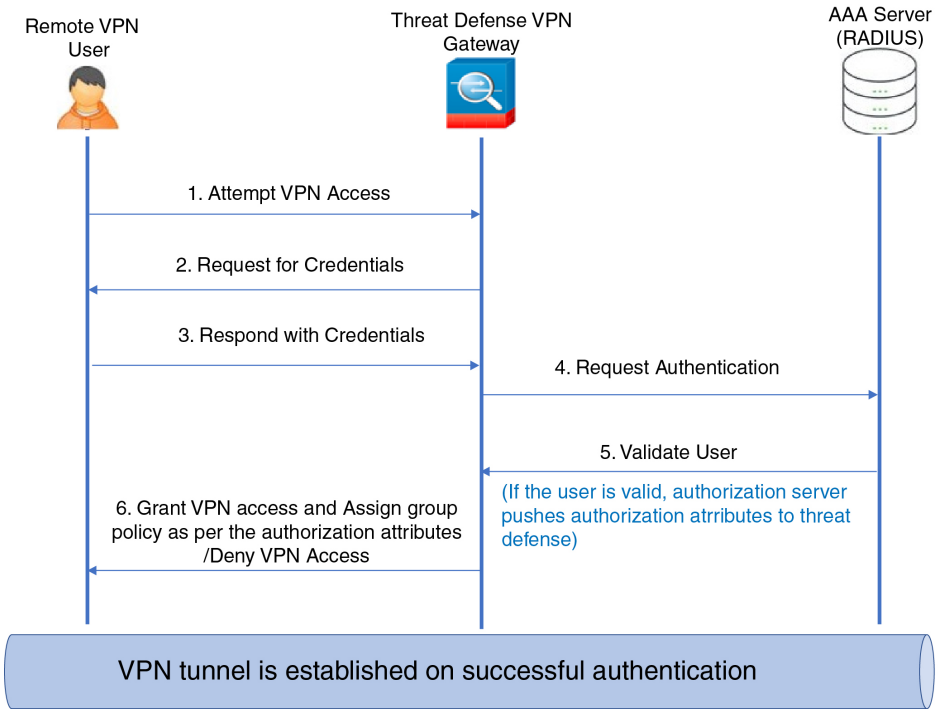

Related Topics

Configure Group Policy Objects, on page 1466
Configure Connection Profile Settings, on page 1594

Override the Selection of Group Policy or Other Attributes by the Authorization Server

When a remote access VPN user connects to the VPN, the group policy and other attributes configured in the connection profile are assigned to the user. However, the remote access VPN system administrator can delegate the selection of group policy and other attributes to the authorization server by configuring ISE or the RADIUS Server to set the Authorization Profile for a user or user-group. Once users are authenticated, these specific authorization attributes are pushed to the Firewall Threat Defense device.

Before you begin

Ensure that you configure a remote access VPN policy with RADIUS as the authentication server.

Procedure

Step 1 On your Secure Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select a remote access policy and click **Edit**.

Step 3 Select RADIUS or ISE as the authorization server if not configured already.

Step 4 Select **AdvancedGroup Policies** and add the required group policy. For detailed information about a group policy object, see Configure Group Policy Objects, on page 1466.

You can map only one group policy to a connection profile; but you can create multiple group policies in a remote access VPN policy. These group policies can be referenced in ISE or the RADIUS server and configured to override the group policy configured in the connection profile by assigning the authorization attributes in the authorization server.

Step 5 Deploy the configuration on the target Firewall Threat Defense device.

Step 6 On the authorization server, create an Authorization Profile with RADIUS attributes for IP address and downloadable ACLs.

When the group policy is configured in the authorization server selected for remote access VPN, the group policy overrides the group policy configured in the connection profile for the remote access VPN user after the user is authenticated.

Related Topics

Configure Group Policy Objects, on page 1466

Deny VPN Access to a User Group

When you do not want an authenticated user or user group to be able to use VPN, you can configure a group policy to deny VPN access. You can configure a group policy in a remote access VPN policy and reference it in the ISE or RADIUS server configuration for authorization.

Before you begin

Ensure that you have configured remote access VPN using the Remote Access Policy wizard and configured authentication settings for the remote access VPN policy.

Procedure

Step 1 On your Secure Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select a remote access policy and click **Edit**.

Step 3 Select **AdvancedGroup Policies**.

Step 4 Select a group policy and click **Edit** or add a new group policy.

Step 5 Select **AdvancedSession Settings** and set **Simultaneous Login Per User** to 0 (zero).
This stops the user or user group from connecting to the VPN even once.

Step 6 Click **Save** to save the group policy and then save the remote access VPN configuration.

Step 7 Configure ISE or the RADIUS server to set the Authorization Profile for that user/user-group to send IETF RADIUS Attribute 25 and map to the corresponding group policy name.

Step 8 Configure the ISE or RADIUS server as the authorization server in the remote access VPN policy.

Step 9 Save and deploy the remote access VPN policy.

Related Topics

Configure Connection Profile Settings, on page 1594

Restrict Connection Profile Selection for a User Group

When you want to enforce a single connection profile on a user or user group, you can choose to disable the connection profile so that the group alias or URLs are not available for the users to select when they connect using the AnyConnect VPN module of Cisco Secure Client.

For example, if your organization wants to use specific configurations for different VPN user groups such as mobile users, corporate-issued laptop users, or personal laptop users, you can configure connection a profile specific to each of these user groups and apply the appropriate connection profile when the user connects to the VPN.

The AnyConnect VPN module of Cisco Secure Client, by default, shows a list of the connection profiles (by connection profile name, alias, or alias URL) configured in Firewall Management Center and deployed on Firewall Threat Defense. If custom connection profiles are not configured, AnyConnect VPN module of Cisco Secure Client shows the *DefaultWEBVPNGroup* connection profile. Use the following procedure to enforce a single connection profile for a user group.

Before you begin

- On your Secure Firewall Management Center web interface, configure remote access VPN using the remote access VPN policy wizard with Authentication Method as 'Client Certificate Only' or 'Client Certificate + AAA'. Choose the username fields from the certificate.
- Configure ISE or RADIUS server for authorization and associate the group policy with the authorization server.

Procedure

Step 1 On your Secure Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select a remote access policy and click **Edit**.

Step 3 Select **Access Interfaces** and disable **Allow users to select connection profile while logging in**.

Step 4 Click **Advanced Certificate Maps**.

Step 5 Select **Use the configured rules to match a certificate to a Connection Profile**.

Step 6 Select the **Certificate Map Name** or click the **Add** icon to add a certificate rule.

Step 7 Select the **Connection Profile**, and click **Ok**.
With this configuration, when a user connects from the AnyConnect VPN module of Cisco Secure Client, the user will have the mapped connection profile and will be authenticated to use the VPN.

Related Topics

Configure Group Policy Objects, on page 1466
Configure Connection Profile Settings, on page 1594

Update the Secure Client Profile for Remote Access VPN Clients

Secure Client Profile is an XML file that contains an administrator-defined end user requirements and authentication policies to be deployed on a VPN client system as part of Secure Client. It makes the preconfigured network profiles available to end users.

You can use the GUI-based Secure Client Profile Editor, an independent configuration tool, to create them. The standalone profile editor can be used to create a new or modify existing Secure Client profile. You can download the profile editor from Cisco Software Download Center.

See the Secure Client Profile Editor chapter in the appropriate release of the Cisco Secure Client (including AnyConnect) Administrator Guide for details.

Before you begin

- Ensure that you have configured remote access VPN using the Remote Access Policy wizard and deployed the configuration on Firewall Threat Defense device. See Create a New Remote Access VPN Policy, on page 1584.
- On your Secure Firewall Management Center web interface, go to **Objects** > **Object Management** > **VPN** > **Secure Client File** and add the new Secure Client image.

Procedure

Step 1 On your Secure Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**.

Step 2 Select a remote access VPN policy and click **Edit**.

Step 3 Select the connection profile that includes the client profile to be edited, and click **Edit**.

Step 4 Click **Edit Group Policy** > **Secure Client** > **Profiles**.

Step 5 Select the client profile XML file from the list or click **Add** to add a new client profile.

Step 6 Save the group policy, connection profile, and then the remote access VPN policy.

Step 7 Deploy the changes.
Changes to the client profile will be updated on the VPN clients when they connect to the remote access VPN gateway.

Related Topics

Configure Group Policy Objects, on page 1466

RADIUS Dynamic Authorization

Secure Firewall Threat Defense has the capability to use RADIUS servers for user authorization of VPN remote access and firewall cut-through-proxy sessions using dynamic access control lists (ACLs) or ACL names per user. To implement dynamic ACLs for dynamic authorization or RADIUS Change of Authorization (RADIUS CoA), you must configure the RADIUS server to support them. When the user tries to authenticate, the RADIUS server sends a downloadable ACL or ACL name to the Firewall Threat Defense. Access to a

given service is either permitted or denied by the ACL. Secure Firewall Threat Defense deletes the ACL when the authentication session expires.

Related Topics

Add a RADIUS Server Group, on page 1364
Interface, on page 1398
Configuring RADIUS Dynamic Authorization, on page 1659
RADIUS Server Attributes for Secure Firewall Threat Defense, on page 1602

Configuring RADIUS Dynamic Authorization

Before you begin:

- Only one interface can be configured in the security zone or interface group if it is referred in a RADIUS Server.
- A dynamic authorization enabled RADIUS server requires Secure Firewall Threat Defense 6.3 or later for the dynamic authorization to work.
- Interface selection in RADIUS server is not supported on Secure Firewall Threat Defense 6.2.3 or earlier versions. The interface option will be ignored during deployment.
- Firewall Threat Defense posture VPN does not support group policy change through dynamic authorization or RADIUS change of authorization (CoA).

Table 88: Procedure

	Do This	More Info
Step 1	Log on to your Secure Firewall Management Center web interface.	
Step 2	Configure a RADIUS server object with dynamic authorization.	RADIUS Server Group Options, on page 1364
Step 3	Configure a route to ISE server through an interface enabled for change of authorization (CoA) to establish connectivity from Firewall Threat Defense to RADIUS server through routing or a specific interface.	RADIUS Server Group Options, on page 1364 Ways to Configure the Cisco Identity Services Engine (Cisco ISE) Identity Source, on page 2512
Step 4	Configure a remote access VPN policy and select the RADIUS server group object that you have created with dynamic authorization.	Create a New Remote Access VPN Policy, on page 1584
Step 5	Configure the DNS server details and domain-lookup interfaces using the Platform Settings.	Configure DNS, on page 1589 DNS Server Group, on page 1384

	Do This	More Info
Step 6	Configure a split-tunnel in group policy to allow DNS traffic through Remote Access VPN tunnel if the DNS server is reachable through VNP network.	Configure Group Policy Objects, on page 1466
Step 7	Deploy the configuration changes.	Deploy Configuration Changes, on page 259

Two-Factor Authentication

You can configure two-factor authentication for the remote access VPN. With two-factor authentication, the user must supply a username and static password, plus an additional item such as an RSA token or a passcode. Two-factor authentication differs from using a second authentication source in that two-factor is configured on a single authentication source, with the relationship to the RSA server tied to the primary authentication source.

Secure Firewall Threat Defense supports RSA tokens and Duo Push authentication requests to Duo Mobile for the second factor in conjunction with any RADIUS or AD server as the first factor in the two-factor authentication process.

Configuring RSA Two-Factor Authentication

About this task:

You can configure the RADIUS or AD server as the authentication agent in the RSA server, and use the server in Secure Firewall Management Center as the primary authentication source in the remote access VPN.

When using this approach, the user must authenticate using a username that is configured in the RADIUS or AD server, and concatenate the password with the one-time temporary RSA token, separating the password and token with a comma: *password*,*token*.

In this configuration, it is typical to use a separate RADIUS server (such as one supplied in Cisco ISE) to provide authorization services. You would configure the second RADIUS server as the authorization and, optionally, accounting server.

Before you begin:

Ensure that the following configurations are complete before configuring RADIUS two-factor authentication on Secure Firewall Threat Defense:

On the RSA Server

- Configure RADIUS or Active Directory server as an authentication agent.
- Generate and download the configuration (*sdconf.rec*) file.
- Create a token profile, assign the token to the user, and distribute the token to the user. Download and install the token on the remote access VPN client system.

For more information, see RSA SecureID Suite documentation.

On the ISE Server

- Import the configuration (*sdconf.rec*) file generated on the RSA server.
- Add the RSA server as the external identity source and specify the shared secret.

Table 89: Procedure

	Do This	More Info
Step 1	Log on to your Secure Firewall Management Center web interface.	
Step 2	Create a RADIUS server group.	RADIUS Server Group Options, on page 1364
Step 3	Create a RADIUS Server object within the new RADIUS server group, with RADIUS or AD server as the host and with a timeout of 60 seconds or more.	RADIUS Server Group Options, on page 1364 **Note** The RADIUS or AD server must be the same server that is configured as the authentication agent in RSA server. For two-factor authentication, make sure that the timeout is updated to 60 seconds or more in the Secure Client Profile XML file as well.
Step 4	Configure a new remote access VPN policy using the wizard or edit an existing remote access VPN policy.	Create a New Remote Access VPN Policy, on page 1584
Step 5	Select RADIUS as the authentication server and then select the newly-created RADIUS server group as the authentication server.	Configure AAA Settings for Remote Access VPN, on page 1597
Step 7	Deploy the configuration changes.	Deploy Configuration Changes, on page 259

Configuring Duo Two-Factor Authentication

About this task:

You can configure the Duo RADIUS server as the primary authentication source. This approach uses the Duo RADIUS Authentication Proxy. (You cannot use a direct connection with the Duo Cloud Service over LDAPS.)

For the detailed steps to configure Duo, see https://duo.com/docs/cisco-firepower.

You would then configure Duo to forward authentication requests directed to the proxy server to use another RADIUS server, or an AD server, as the first authentication factor, and the Duo Cloud Service as the second factor.

When using this approach, the user must authenticate using a username that is configured on both the Duo Cloud or web server, and the associated RADIUS server. The user must enter the password configured in the RADIUS server, followed by one of the following Duo codes:

- **Duo-passcode**. For example, *my-password*,123456.

- **push**. For example, *my-password*,**push**. Use push to tell Duo to send a push authentication to the Duo Mobile app, which the user must have already installed and registered.
- **sms**. For example, *my-password*,**sms**. Use **sms** to tell Duo to send an SMS message with a new batch of passcodes to the user's mobile device. The user's authentication attempt will fail when using **sms**. The user must then re-authenticate and enter the new passcode as the secondary factor.
- **phone**. For example, *my-password*,**phone**. Use **phone** to authenticate using phone callback.

For more information on login options with examples, see https://guide.duo.com/anyconnect.

Before you begin:

Before configuring two-factor authentication with Duo Authentication Proxy on Firewall Threat Defense, ensure that you complete the following configurations:

- Configure a working primary authentication (RADIUS or AD) for your remote access VPN users before you begin to deploy Duo.
- Install Duo proxy service on a Windows or Linux machine within your network to integrate Duo with Secure Firewall Threat Defense remote access VPN. This Duo proxy server also acts as a RADIUS server.

  Download and install the most recent Duo authentication proxy from the following location:

  - **Windows**: https://dl.duosecurity.com/duoauthproxy-latest.exe
  - **Linux**: https://dl.duosecurity.com/duoauthproxy-latest-src.tgz
  - Verify the checksum at https://duo.com/docs/checksums#duo-authentication-proxy.

- Configure Duo authentication file `authproxy.cfg`. Follow instructions on the https://duo.com/docs/cisco-firepower#configure-the-proxy page to configure the authentication configuration settings.

  The `authproxy.cfg` configuration file must contain the details for RADIUS or ISE server, Firewall Threat Defense device, Duo proxy server details, Integration Key, Secret key, and API host details.
- Ensure that you have the right API host information in the `authproxy.cfg` file.
- Configure other required settings such as secondary authentication factor in the newly installed Duo proxy server at **Duo Security Server** > **Duo Admin Panel** > **Applications** > **CISCO RADIUS VPN**.

Table 90: Procedure

	Do This	More Info
Step 1	Log on to your Secure Firewall Management Center web interface.	
Step 2	Create a RADIUS server group.	RADIUS Server Group Options, on page 1364

	Do This	More Info
Step 3	Create a RADIUS Server object within the new RADIUS server group with Duo proxy server as the host with a timeout of 60 seconds or more.	RADIUS Server Options, on page 1366 **Note** For two-factor authentication, make sure that the timeout is updated to 60 seconds or more in the Secure Client Profile XML file as well.
Step 4	Configure a new remote access VPN policy using the wizard or edit an existing remote access VPN policy.	Create a New Remote Access VPN Policy, on page 1584
Step 5	Select RADIUS as the authentication server and then select the RADIUS server group created with the Duo proxy server as the authentication server.	Configure AAA Settings for Remote Access VPN, on page 1597
Step 7	Deploy the configuration changes.	Deploy Configuration Changes, on page 259

Secondary Authentication

Secondary authentication or double authentication in Secure Firewall Threat Defense adds an additional layer of security to remote access VPN connections by using two different authentication servers. With secondary authentication enabled, the Secure Client VPN users must provide two sets of credentials to login to the VPN gateway.

Secure Firewall Threat Defense remote access VPN supports secondary authentication in AAA Only and Client Certificate & AAA authentication methods.

Figure 436: Remote Access VPN Secondary or Double Authentication

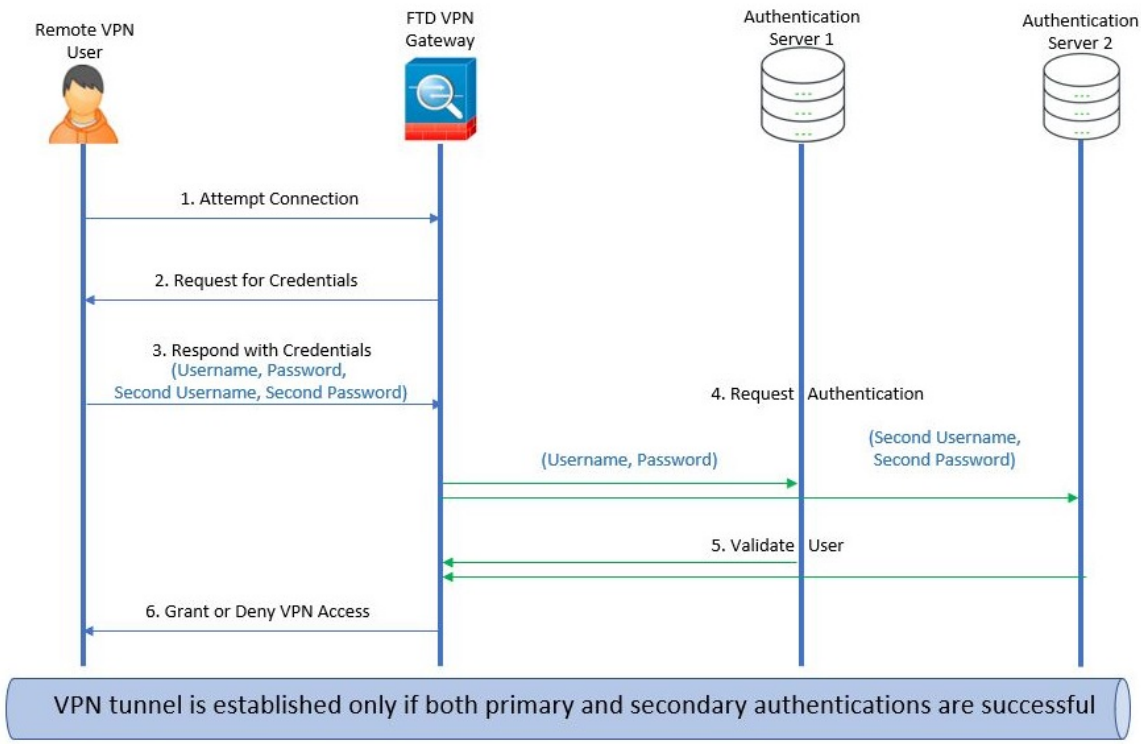

Related Topics

Configure Remote Access VPN Secondary Authentication, on page 1664

Configure Remote Access VPN Secondary Authentication

When remote access VPN authentication is configured to use both client certificate and authentication sever, VPN client authentication is done using both the client certificate validation and AAA server.

Before you begin

- Configure two authentication (AAA) servers— the primary and secondary authentication servers, and required identity certificates. The authentication servers can be RADIUS server, and AD or LDAP realms.
- Ensure that the AAA servers are reachable from the Secure Firewall Threat Defense device for the remote access VPN configuration to work. Configure routing (at **Devices > Device Management > Edit Device > Routing**) to ensure connectivity to the AAA servers.

Procedure

Step 1 On your Secure Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access** .

Step 2 Select a remote access policy and click **Edit**; or click **Add** to create a new remote access VPN policy.

Step 3 For a new remote access VPN policy, configure the authentication while selecting connection profile settings. For an existing configuration, select the connection profile that includes the client profile, and click **Edit**.

Step 4 Click **AAA** > **Authentication Method**, **AAA** or **Client Certificate & AAA**.

- When you select the **Authentication Method** as:

  Client Certificate & AAA—Authentication is done using both client certificate and AAA server.

  - **AAA**—If you select the **Authentication Server** as **RADIUS**, by default, the Authorization Server has the same value. Select the **Accounting Server** from the drop-down list. Whenever you select **AD** and **LDAP** from the Authentication Server drop-down list, you must manually select the **Authorization Server** and **Accounting Server** respectively.
  - Whichever authentication method you choose, select or deselect **Allow connection only if user exists in authorization database**.

- **Use secondary authentication** — Secondary authentication is configured in addition to primary authentication to provide additional security for VPN sessions. Secondary authentication is applicable only to **AAA only** and **Client Certificate & AAA** authentication methods.

  Secondary authentication is an optional feature that requires a VPN user to enter two sets of username and password on the Secure Client login screen. You can also configure to pre-fill the secondary username from the authentication server or client certificate. Remote access VPN authentication is granted only if both primary and secondary authentications are successful. VPN authentication is denied if any one of the authentication servers is not reachable or one authentication fails.

  You must configure a secondary authentication server group (AAA server) for the second username and password before configuring secondary authentication. For example, you can set the primary authentication server to an LDAP or Active Directory realm and the secondary authentication to a RADIUS server.

  Note
  By default, secondary authentication is not required.

  Authentication Server— Secondary authentication server to provide secondary username and password for VPN users.

  Select the following under **Username for secondary authentication**:

  - **Prompt**: Prompts the users to enter the username and password while logging on to VPN gateway.
  - **Use primary authentication username**: The username is taken from the primary authentication server for both primary and secondary authentication; you must enter two passwords.
  - **Map username from client certificate**: Prefills the secondary username from the client certificate.
    - If you select **Map specific field** option, which includes the username from the client certificate. The **Primary** and **Secondary** fields display default values: **CN (Common Name)** and **OU (Organisational Unit)** respectively. If you select the **Use entire DN (Distinguished Name) as username** option, the system automatically retrieves the user identity.

      See **Authentication Method** descriptions for more information about primary and secondary field mapping.
    - **Prefill username from certificate on user login window**: Prefills the secondary username from the client certificate when the user connects via Secure Client.
      - **Hide username in login window**: The secondary username is pre-filled from the client certificate, but hidden to the user so that the user does not modify the pre-filled username.

- **Use secondary username for VPN session**: The secondary username is used for reporting user activity during a VPN session.

For more information, see Configure AAA Settings for Remote Access VPN, on page 1597.

Related Topics

Configure Connection Profile Settings, on page 1594

Single Sign-On Authentication with SAML 2.0

About SAML Single Sign-On Authentication

Security Assertion Markup Language (SAML) is an open standard for logging users into applications using their sessions in another context. Organizations already know the identity of users when users log in to their Active Directory (AD) domain or the intranet. They use this identity information to log in users to other applications, such as web-based applications using SAML. Individual applications do not need to store credentials and users do not have to remember and manage different sets of credentials for individual applications. SAML single sign-on (SSO) works by transferring the user's identity from one place (the identity provider) to another (the service provider).

SAML Single Sign-On with Secure Firewall Threat Defense

The Secure Firewall Threat Defense device supports SAML 2.0 single sign-on (SSO) authentication for remote access VPN connections using the Secure Client. You need the following to configure SAML 2.0 SSO on Secure Firewall Threat Defense:

- **Identity Provider (IdP)**—The Duo Access Gateway acts as the identity provider to perform user authentication and issues assertions.
- **Service Provider (SP)**—The Firewall Threat Defense device acts as the service provider and obtains the authentication assertion from the identity provider.
- **VPN Client**—The Secure Client performs SAML 2.0 authentication via the embedded browser.

Guidelines and Limitations for SAML 2.0

- Firewall Threat Defense supports the following signatures for SAML authentication:
  - SHA1 with RSA and HMAC
  - SHA2 with RSA and HMAC
- Firewall Threat Defense supports SAML 2.0 Redirect-POST binding , which is supported by all SAML IdPs.
- The Firewall Threat Defense functions as a SAML SP only. It cannot act as an Identity Provider in gateway mode or peer mode.
- You can enforce an access policy on a SAML-authenticated user if you have an associated identity policy with an AD realm matching the SAML domain. However, it does not work for Azure AD SAML because it requires additional mapping from the tenant ID of the Azure AD to an associated realm ID on the threat defense device.

- Having SAML authentication attributes available in DAP evaluation (similar to RADIUS attributes sent in RADIUS auth response from AAA server) is not supported. Firewall Threat Defense supports SAML enabled group policy on DAP policy; however, you cannot check the username attribute while using SAML authentication, because the username attribute is masked by the SAML Identity provider.
- Firewall Threat Defense administrators need to ensure clock synchronization between the Firewall Threat Defense and the SAML IdP for proper handling of authentication assertions and proper timeout behavior.
- Firewall Threat Defense administrators have the responsibility to maintain a valid signing certificate on both Firewall Threat Defense and IdP considering the following:
  - The IdP signing certificate is mandatory when configuring an IdP on the Firewall Threat Defense.
  - The Firewall Threat Defense does not do a revocation check on the signing certificate received from the IdP.
- In SAML assertions, there are NotBefore and NotOnOrAfter conditions. The Firewall Threat Defense SAML configured timeout interacts with these conditions as follows:
  - Timeout overrides NotOnOrAfter if the sum of NotBefore and timeout is earlier than NotOnOrAfter.
  - If NotBefore + timeout is later than NotOnOrAfter, then NotOnOrAfter takes effect.
  - If the NotBefore attribute is absent, the Firewall Threat Defense denies the login request. If the NotOnOrAfter attribute is absent and SAML timeout is not set, Firewall Threat Defense denies the login request.
- Firewall Threat Defense does not work with Duo in a deployment using an internal SAML, which forces the Firewall Threat Defense to proxy for the client to authenticate, due to the FQDN change that occurs during challenge/response for Two-factor authentication (push, code, password).
- When using SAML with Secure Client, follow these guidelines:
  - Untrusted server certificates are not allowed in the embedded browser.
  - The embedded browser SAML integration is not supported in CLI or SBL modes.
  - SAML authentication established in a web browser is not shared with Secure Client and vice versa.
  - Depending on the configuration, various methods are used when connecting to the headend with the embedded browser. For example, while Secure Client might prefer an IPv4 connection over an IPv6 connection, the embedded browser might prefer IPv6, or vice versa. Similarly, Secure Client may fall back to no proxy after trying proxy and getting a failure, while the embedded browser may stop navigation after trying proxy and getting a failure.
  - You must synchronize your Firewall Threat Defense's Network Time Protocol (NTP) server with the IdP NTP server in order to use the SAML feature.
  - You cannot access internal servers with SSO after logging in using an internal IdP.
  - The SAML IdP NameID attribute determines the user's username and is used for authorization, accounting, and VPN session database.
- SAML does not support Start Before Logon (SBL).
- Multiple attributes received with a SAML assertion are not supported.

- A firewall uses the Identity Provider Entity ID as the SAML object name from the single sign-on server object created in the Firewall Management Center (**Objects** > **Object Management** > **AAA Server** > **Single Sign-on Server**). Hence, you cannot use multiple SAML objects with the same Identity Provider Entity ID on a single firewall.

Configuring a SAML Single Sign-On Authentication

Before you begin

Ensure that you have done the following before you configure SAML single sign-on with Firewall Threat Defense remote access VPN:

- Create an account with Duo.
- Download and install the Duo Access Gateway.
- Obtain the following from your SAML identity provider (Duo).
  - Identity Provider Entity ID URL
  - Sign-in URL
  - Sign-out URL
  - Identity provider certificate
- Create a SAML single sign-on server object. For more information, see Add a Single Sign-on Server, on page 1367.

Note You can create a single sign-on server object in the **Connection Profile** settings when you create a new policy using the Remote Access VPN policy Wizard.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** next to the remote access VPN policy for which you want to configure SAML authentication. If you want to create a new policy, click **Add**.

Step 3 Click **Edit** on the connection profile that you want to modify.

Step 4 Choose **AAA** settings and select **SAML** from the **Authentication Method** drop-down.

Step 5 Choose the required SAML single sign-on server as the **Authentication Server**.

Step 6 Configure the required settings for the remote access VPN.

Step 7 Save and deploy the remote access VPN policy on your Firewall Threat Defense device.

Related Topics

Configure AAA Settings for Remote Access VPN, on page 1597

Configuring SAML Authorization

About SAML Authorization

SAML authorization supports user attributes delivered in SAML assertions within the AAA and Dynamic Access Policy (DAP) frameworks. You can configure the SAML assertion attributes on the Identity Provider as name-value pairs which then parses as strings. The attributes received are made available to DAP so that they can be used when defining selection criteria within a DAP record. The SAML assertion *cisco_group_policy* is used to determine the Group Policy to be applied to the VPN session.

Dynamic Access Policy Attribute Representation

In the DAP table, the DAP attributes are represented in the following format:

```
aaa.saml.name = "value”
```

Example, *aaa.saml.department = ”finance"*

This attribute can be used in DAP selection as follows:

```
<attr>
<name>aaa.saml.department</name>
<value>finance</value>
<operation>EQ</operation>
</attr>
```

Multi-Valued Attributes

Multi-valued attributes are also supported in DAP and the DAP table is indexed :

```
aaa.saml.name.1 = "value”
aaa.saml.name.2 = "value"
```

Active Directory memberOf Attributes

The Active Directory (AD) memberOf attribute receives a special processing that is consistent with the way it is handled through an LDAP query.

Group names are represented by the CN attribute of the DN.

Example Attributes received from the authorization server:

```
memberOf = "CN=FTD-VPN-Group,OU=Users,OU=TechspotUsers,DC=techspot,DC=us"
 memberOf = "CN=Domain Admins,OU=Users,DC=techspot,DC=us”
```

Dynamic Access Policy attributes:

```
aaa.saml.memberOf.1 = "FTD-VPN-Group”
aaa.saml.memberOf.2 = "Domain Admins"
```

Interpretation of the cisco_group_policy Attribute

A group-policy can be specified by a SAML assertion attribute. When an attribute "cisco_group_policy" is received by the Firewall Threat Defense, the corresponding value is used to select the connection group-policy

Configure SAML Authorization

Before you begin

Ensure that you have configured a single sign-on server like DUO and completed the required Identity Provider (IdP) and Service Provider (SP) settings.

For more information, see Single Sign-On Authentication with SAML 2.0, on page 1666.

Procedure

Step 1 Configure a single sign-on server object if not configured already.

a) Choose **Object** > **Object Management** > **AAA Server** > **Single Sign-on Server**
b) Click **Add Single Sign-on Server**.
c) Enter the single sign-on server details and click **Save**.

For more information, see Add a Single Sign-on Server, on page 1367.

Step 2 Configure SAML authentication in the remote access VPN connection profile.

a) Choose **Devices** > **Remote Access**.
b) Click **Edit** on the remote access VPN policy for which you want to configure SAML authorization or create a new policy.
c) Edit the required connection profile and select **AAA**.
d) Select the single sign-on server object from the **Authentication Server** drop-down.
e) Save the remote access VPN configuration.

Step 3 Match a SAML criteria in DAP policy.

a) Select **Devices** > **Dynamic Access Policy**.
b) Create a new DAP or edit an existing one.
c) Create a DAP record or edit and existing record.
d) Click **AAA Criteria** > **SAML Criteria** > **Add SAML Criteria**.
e) Create a SAML criteria based on the SAML assertions returned by the SSO server.

Step 4 Deploy the remote access VPN configuration.

Related Topics

Configure Connection Profile Settings, on page 1594
Firewall Threat Defense Group Policy Objects, on page 1465

Advanced Secure Client Configurations

Configure Secure Client Modules on a Firewall Threat Defense

Secure Client can integrate with various Cisco endpoint security solutions and offer enhanced security using different Secure Client modules.

You can use the managed headend Firewall Threat Defense to distribute and manage Secure Client modules to the endpoints. When a user connects to the Firewall Threat Defense, it downloads and installs Secure Client and the required modules on the endpoint.

In version 6.7 and later, you can use the headend Firewall Threat Defense, managed by a Firewall Management Center, to distribute and manage Secure Client modules to the endpoints. These modules then integrate with the corresponding Cisco endpoint security solution.

In versions 6.4 to 6.6, you can enable these modules and profiles on a Firewall Threat Defense using FlexConfig. For more information, see Configure AnyConnect Modules and Profiles Using FlexConfig.

Benefits

If you use a Firewall Threat Defense to distribute and manage Secure Client modules to the endpoints, you can easily perform the following tasks:

- Distribute and manage Secure Client modules and profiles on each endpoint.
- Upgrade Secure Client on each endpoint.

Types of Secure Client Modules

AMP Enabler

Use this module to deploy Secure Endpoint, on endpoints. The module pushes Secure Endpoint to endpoints from a server hosted locally within the enterprise. This module provides an additional security agent that detects potential malware threats in the network, removes these threats, and protects the enterprise.

In Cisco Secure Client 5.0, AMP Enabler is only for macOS. Cisco Secure Client for Windows offers full integration with Cisco Secure Endpoint.

ISE Posture

Use this module to perform endpoint posture checks such as antivirus, antispyware, operating system and so on using Cisco Identity Services Engine (ISE) and assess the endpoint's compliance. ISE provides next generation identity and access control policy. ISE Posture performs a client-side evaluation. The client receives the posture requirement policy from the headend, performs the posture data collection, compares the results against the policy, and sends the assessment results back to the headend.

Network Visibility

Use this module to monitor the endpoint application usage using the Network visibility module. You can uncover potential behavior anomalies and make informed network design decisions. It enhances the enterprise administrator's ability to do capacity and service planning, auditing, compliance, and security analytics. You can share the usage data with NetFlow analysis tools such as Cisco Stealthwatch.

Umbrella Roaming Security

Use this module for a DNS-layer security using the Cisco Umbrella Roaming Security service. Cisco Umbrella provides content filtering, multiple policies, robust reporting, active directory integration, and much more.

Web Security

Use this module to enable Cisco Secure Web Appliance (SWA), powered by Cisco Talos. This module protects the endpoint by blocking risky sites and testing unknown sites before allowing users to access them. It can deploy web security either through the on-prem WSA or the cloud-based Cisco Cloud Web Security. This module is not part of the AnyConnect package from release 4.5 and in Secure Client 5.0.

Network Access Manager

This module provides a secure layer 2 network and performs device authentication to access wired and wireless networks. Network Access Manager manages user and device identity and the network access protocols required for secure access.

Network Access Manager is not supported on macOS or Linux.

Start Before Login

Start Before Login (SBL) allows users to establish their VPN connection to the enterprise infrastructure before logging onto Windows. After the SBL module installation, you must enable SBL in the Secure Client VPN profile and add it to the remote access VPN group policy.

DART

Diagnostics and Reporting Tool (DART) collates system logs and other diagnostic information to troubleshoot AnyConnect installation and connection problems. You can send this data to Cisco TAC for troubleshooting.

By default, DART is not enabled in new RA VPN group policies for 6.7 and later versions. In 6.6 and earlier versions, DART is enabled by default.

Feedback

The customer experience feedback (CEF) module provides information about which features and modules you use and have enabled. This information gives an insight into the user experience so that Cisco can continue to improve the quality, reliability, performance, and user experience of the Cisco Secure Client. Secure Client does not download the Feedback module to the endpoint. The feedback data is sent to the Cisco Feedback Server.

Prerequisites for Configuring Secure Client Modules

- Configure the associated products depending on the module that you are going to use.
- Download the following Secure Client-related packages from Cisco Software Download Center to your local host.
  - Cisco Secure Client Headend Deployment Package for the required platforms.

    This package is for the headend and contains all the Secure Client modules. For Windows, the filename is cisco-secure-client-win-5.0.03076-webdeploy-k9.pkg.
  - Profile Editor: Create profiles for the modules that require profiles.

    Secure Client needs a Secure Client profile for some of the modules. A profile contains configurations to enable the modules and connect to the corresponding security services. The profile editor supports only Windows.

    The following table lists if the modules require a client profile:

Secure Client Module	Requires a Client Profile
AMP Enabler	Yes
ISE Posture	Yes
Network Access Manager	Yes
Network Visibility Module	Yes
Umbrella Roaming Secure Module	Yes
Feedback	Yes
Web Security	Yes
DART	No
Start Before Login	No

- Licensing
  - You need one of the following Secure Client licenses: Secure Client Premier, Secure Client Advantage, or Secure Client VPN Only
  - Your Firewall Management Center Essentials license must allow export-controlled functionality.

    Choose **System > Licenses > Smart Licenses** to verify this functionality in the management center.

Guidelines for Configuring Secure Client Modules

- All Secure Client modules are supported from AnyConnect 4.8 and later, and Secure Client 5.0.
- Different modules support profiles with different file extensions. The following table lists the modules and the supported file extensions of their profiles:

Table 91: Supported File Extensions of Profiles

Module	File Extension
AMP Enabler	*.xml, *.asp
Feedback	*.xml
ISE Posture	*.xml, *.isp
Network Access Manager	*.xml, *.nsp
Network Visibility	*.xml, *.nvmsp
Umbrella Roaming Security	*.xml, *.json
Web Security	*.xml, *.wsp, *.wso

- You can add only one entry per client module. You can edit or delete an entry for a module.

- If you plan to use ISE posture and Network Access Manager modules on a Windows OS, you must install Network Access Manager before you use the ISE Posture module.
- If you enable the Umbrella Roaming Security module, ensure that you disable the **Always send DNS requests over tunnel** option under split tunnelling in the VPN group policy.
- If you want to use SBL, then you must enable SBL in the Secure Client VPN profile.

Install Secure Client Modules using a Firewall Threat Defense

Before you begin

Ensure that you review the Prerequisites for Configuring Secure Client Modules, on page 1672 and Guidelines for Configuring Secure Client Modules, on page 1673 topics.

Procedure

Step 1 The administrator creates profiles, if needed, for the required Secure Client modules.

Step 2 The administrator uses the Firewall Management Center to:

a) Configure the modules and add the profiles in the remote access VPN group policy.

b) Deploy the configuration on the Firewall Threat Defense.

Step 3 The user uses Secure Client to initiate a VPN connection to the Firewall Threat Defense.

Step 4 The Firewall Threat Defense authenticates the user.

Step 5 The Secure Client checks for updates.

Step 6 The Firewall Threat Defense distributes the Secure Client modules and the profiles on the endpoint.

What to do next

Configure a Remote Access VPN Group Policy with Secure Client Modules, on page 1674.

Configure a Remote Access VPN Group Policy with Secure Client Modules

To install and update the Secure Client modules on the endpoint using a Firewall Threat Defense managed by a Firewall Management Center, you must update the remote access VPN group policy with the Secure Client module configurations.

Before you begin

Ensure that you have configured a remote access VPN policy in the Firewall Management Center.

Procedure

Step 1 Choose **Devices > Remote Access**.

Step 2 Select a remote access VPN policy and click **Edit**.

Step 3 Select a connection profile and click **Edit**.

Step 4 Click **Edit Group Policy**.

Step 5 Click the **Secure Client** tab.

Step 6 Click **Client Modules**.

Step 7 Click +.

Step 8 Choose a module from the **Client Module** drop-down list.

Step 9 Choose a profile for the module from the **Profile to download** drop-down list or click + to add a profile.

Step 10 Check the **Enable module download** check box to download the module on the endpoint.

Step 11 Click **Add**.

Step 12 Repeat steps 7 to 11 if you want to add more modules.

Step 13 Click **Save**.

What to do next

1. Deploy the configuration on the threat defense.
2. Launch the Secure Client, select the VPN profile, and connect to the VPN. Secure Client installs the configured modules on it.
3. Verify the configuration. For more information, see Verify Secure Client Modules Configuration, on page 1675.

Verify Secure Client Modules Configuration

On the Firewall Threat Defense

Use the following commands on the Firewall Threat Defense to view the profiles and the Secure Client modules configuration:

- **show disk0:**- View the profiles and their configuration.
- **show run webvpn** - View details of the Secure Client configurations.
- **show run group-policy <ravpn_group_policy_name>** - View details of the RA VPN group policy for Secure Client.
- **show vpn-sessiondb anyconnect** - View details of the active Secure Client VPN sessions.

On the Endpoint

1. Use the Secure Client to establish a VPN connection to the Firewall Threat Defense.
2. Verify if the configured modules are downloaded and installed as part of the Secure Client.
3. Verify if the configured profiles, if any, are available in the locations documented in Profile Locations for all Operating Systems.

On the Firewall Management Center

You can monitor remote access VPN active sessions on the Firewall Management Center using the Remote Access VPN dashboard (**Overview > Remote Access VPN**). You can quickly determine problems related to user sessions and mitigate the problems for your network and users.

Configure Application-Based (Per App VPN) Remote Access VPN on Mobile Devices

When you use Secure Client to establish a VPN connection from a mobile device, all the traffic including the traffic from personal applications is routed through the VPN.

For mobile devices that run on Android or iOS, you can restrict the applications that use the VPN tunnel. This application-based remote access VPN is called Per App VPN. To use Per App VPN, you must install and configure a third-party Mobile Device Manager (MDM) application. You must define the list of approved applications that can be used over the VPN tunnel in the MDM. You can enable Per App VPN on the Firewall Threat Defense headend so that your MDM can apply your policies on mobile devices.

Benefits

Benefits of restricting the remote access VPN to approved applications include:

- Performance—Limits VPN traffic over the corporate network and frees up resources of the VPN headend.
- Protection—Protects the corporate VPN tunnel from unapproved malicious applications on the mobile device.

Prerequisites and Licensing for Configuring Per App VPN Tunnels

Prerequisites

- Install and configure a third-party Mobile Device Manager (MDM).

  You must configure the applications that will be allowed in the VPN in the MDM itself, not on the Firewall Threat Defense headend device.

- Download the Cisco AnyConnect Enterprise Application Selector from Cisco Software Download Center.

  You need this tool to define the Per App VPN policy.

Licensing

- Secure Client Premier, or Secure Client Advantage.
- Essentials license must allow export-controlled functionality.

  To verify this functionality in the Firewall Management Center, choose **System > Licenses > Smart Licenses**.

Determine the Application IDs for Mobile Applications

Before configuring the Firewall Threat Defense headend to allow application-based VPN from mobile devices, you must determine which applications should be allowed in the tunnel.

We strongly recommend that you configure the per-app policy in the MDM on the user's mobile device. This simplifies the headend configuration. If you decide to configure the list of allowed applications on the headend, you must determine the application IDs for each application on each type of endpoint.

The application ID, called the bundle ID in iOS, is a reverse DNS name. You can use an asterisk as a wildcard. For example, *.* indicates all applications, com.cisco.* indicates all Cisco applications.

To determine the application IDs:

- **Android**—Go to Google Play in a web browser and select the Apps category. Click (or hover over) an application that you want to allow, then look at the URL. The app id is in the URL, on the **id=** parameter. For example, the following URL is for Facebook Messenger, so the app id is com.facebook.orca.

  https://play.google.com/store/apps/details?id=com.facebook.orca

  For applications that are not available through Google Play, such as your own applications, download a package name viewer application to extract the app ID. There are many of these applications available, one of them should provide what you need, but Cisco does not endorse any of them.

- **iOS**—There is no straight-forward way to get the bundle ID. Following is one way to find it:

  1. Use a desktop web browser such as Chrome to search for the application name.
  2. In the search results, look for the link to download the app from the Apple App Store. For example, Facebook Messager would be similar to:

     https://apps.apple.com/us/app/messenger/id454638411
  3. Copy the number after the **id** string. In this example, **454638411**.
  4. Open a new browser window, and add the number to the end of the following URL:

     https://itunes.apple.com/lookup?id=

     For this example: https://itunes.apple.com/lookup?id=454638411
  5. You will be prompted to download a text file, usually named 1.txt. Download the file.
  6. Open the file in a text editor such as WordPad, and search for bundleId. For example:

     "bundleId":"com.facebook.Messenger",

     In this example, the bundle ID is com.facebook.Messenger. Use this as the app ID.

Once you have your list of application IDs, you can configure the policy as explained in .

Configure Application-Based VPN Tunnels

After you install and configure your MDM software, you can enable Per App VPN on the Firewall Threat Defense headend device. Once enabled on the headend, your MDM software will control which applications are tunneled over the VPN to the corporate network.

Before you begin

- Ensure that you have a remote access VPN policy in the Firewall Management Center.
- Configure Per App VPN using MDM and enroll each device to the MDM server.
- Download the Cisco AnyConnect Enterprise Application Selector.

Procedure

Step 1 Use the Cisco AnyConnect Enterprise Application Selector to define the Per App VPN policy.

We recommend that you create a simple **Allow All** policy, and define the allowed applications in the MDM. However, you can specify a list of applications to allow and control the list from the headend. If you want to include specific applications, create a separate rule for each application, using a unique name and the application's app ID. For more information on getting the app IDs, see Determine the Application IDs for Mobile Apps.

To create an **Allow All** policy that supports both Android and iOS platforms using the AnyConnect Enterprise Application Selector:

a) Choose **Android** from the drop-down list as the platform type and configure the following options:

- **Friendly Name**—Enter a name for the policy. For example, Allow_All.
- **App ID**—Enter *.* to match all possible applications.
- Leave the other options.

b) Choose **iOS** from the drop-down list as the platform type and configure the following options:

- **Friendly Name**—Enter a name for the policy. For example, Allow_All.
- **App ID**—Enter *.* to match all possible applications.
- Leave the other options.

c) Choose **Policy > View Policy** to get the base64 encoded string for the policy.

This string contains an encrypted XML file that allows the Firewall Threat Defense to see the policies. Copy this value. You need this string when you configure Per App VPN on the Firewall Threat Defense.

Step 2 Use the Firewall Management Center to enable the Per App on the Firewall Threat Defense headend device.

a) Choose **Devices > Remote Access**.
b) Select a remote access VPN policy and click **Edit**.
c) Select a connection profile and click **Edit**.
d) Click **Edit Group Policy**.
e) Click the **Secure Client** tab.
f) Click **Custom Attributes** and click **+**.
g) Choose **Per App VPN** from the **Secure Client Attribute** drop-down list.
h) Choose an object from the **Custom Attribute Object** drop-down list or click **+** to add an object.

When you add a new custom attribute object for Per App VPN, enter the name, description, and the base64 encoded policy string from the Cisco AnyConnect Enterprise Application Selector.

i) Click **Save**.
j) Click **Add** and click **Save**.

Step 3 Deploy your changes on the Firewall Management Center.

What to do next

1. Launch the Secure Client, select the VPN profile, and connect to the VPN.
2. Verify the configuration. For more information, see Verify Per App Configuration, on page 1679.

Verify Per App Configuration

On the Firewall Threat Defense

Use the following commands on the Firewall Threat Defense to view the Per App configuration:

- **show run webvpn**
- **show run group-policy <ravpn_group_policy_name>**
- **show run anyconnect-custom-data**

On the Endpoint

After the endpoint establishes a VPN connection with the Firewall Threat Defense:

1. Click the **Statistics** icon of the Secure Client.
2. **Tunnel Mode** will be "Application Tunnel" instead of "Tunnel All Traffic."
3. **Tunneled Apps** will list the applications you enabled for tunneling in the MDM.

Remote Access VPN Examples

How to Limit Secure Client Bandwidth Per User

This section provides instructions to limit the maximum bandwidth that the VPN users consume when they connect using the Secure Client to Secure Firewall Threat Defense remote access VPN gateway. You can limit the maximum bandwidth by using a Quality of service (QoS) policy in Firewall Threat Defense, to ensure that a single user or group or users do not take over the entire resource. This configuration lets you give priority to critical traffic, prevent bandwidth hogging, and manage the network. If a When traffic exceeds the maximum rate, the Firewall Threat Defense drops the excess traffic.

Step	Do This	More Info
1	Create and set up a realm.	Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426
2	Create a QoS policy and QoS rule for the user or group available in the newly created realm.	• See Creating a QoS Policy, on page 935 to create a QoS policy. • See Configuring QoS Rules, on page 936 to create a QoS rule.

Step	Do This	More Info
3	Configure remote access VPN policy and select the newly created realm for user authentication.	Create a New Remote Access VPN Policy, on page 1584
4	Deploy the remote access VPN policy.	Deploy Configuration Changes, on page 259

How to Use VPN Identity for User-Id Based Access Control Rules

Step	Do This	More Info
1	Create and set up a realm.	Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.
2	Create an identity policy and add an identity rule.	• See Create an Identity Policy, on page 2555 to create an identity policy. • See Create an Identity Rule, on page 2563 to create an identity rule.
3	Associate the identity policy with an access control policy.	Associating Other Policies with Access Control, on page 1762
4	Configure remote access VPN policy and select the newly created realm for user authentication.	Create a New Remote Access VPN Policy, on page 1584
5	Deploy the remote access VPN policy.	Deploy Configuration Changes, on page 259

Configure Firewall Threat Defense Multiple Certificate Authentication

Multiple Certificate-based Authentication

Multiple certificate-based authentication allows the Firewall Threat Defense to validate the machine or device certificate. Multiple certificates can be enabled for certificate-based authentication in the remote access VPN connection profile. It can be combined with AAA authentication. The multiple certificates option in the remote access VPN connection profile allows certificate authentication of both the machine and user via certificates. This ensures that the device is a corporate-issued device, in addition to authenticating the user's identity certificate to allow RA VPN access. The administrator can choose if the username for the session should be taken from the machine certificate or user certificate.

When multiple certificate-based authentication is configured, two certificates are obtained from the VPN client:

- **First Certificate** —Machine certificate to authenticate the endpoint.
- **Second Certificate**—User certificate to authenticate the VPN user.

For detailed information about Firewall Threat Defense certificates, see Managing Firewall Threat Defense Certificates, on page 1486.

Limitations

- Multiple certificate authentication currently limits the number of certificates to two.
- Secure Client supports only RSA-based certificates.
- Only SHA256, SHA384, and SHA512 based certificates are supported during the Secure Client aggregate authentication.
- Certificate authentication cannot be combined with SAML authentication.

Pre-fill Username from Certificate

The Pre-fill username option allows a field from the certificates to be parsed and used for subsequent AAA authentication (primary and secondary). When two certificates are used for authentication, the Administrator can choose the certificate from which the username should be derived for the prefill functionality. By default, username for prefill is retrieved from the User certificate (second certificate received from Secure Client). The prefilled username is used as the VPN session username when the Certificate Only authentication method is enabled. When AAA and certificate authentication is enabled, VPN session username will be based on the pre-fill option.

Configure Multiple Certificate Authentication for Remote Access VPN

1. On your Secure Firewall Management Center web interface, choose **Devices** > **VPN** > **Remote Access**.
2. Edit an existing remote access policy, or create a new one and then edit.

   See Create a New Remote Access VPN Policy, on page 1584.
3. Select the connection profile to configure multiple certificate authentication, and click **Edit**.

   See Configure Connection Profile Settings, on page 1594.
4. Choose **AAA**, and then select an **Authentication Method**:

Figure 437:

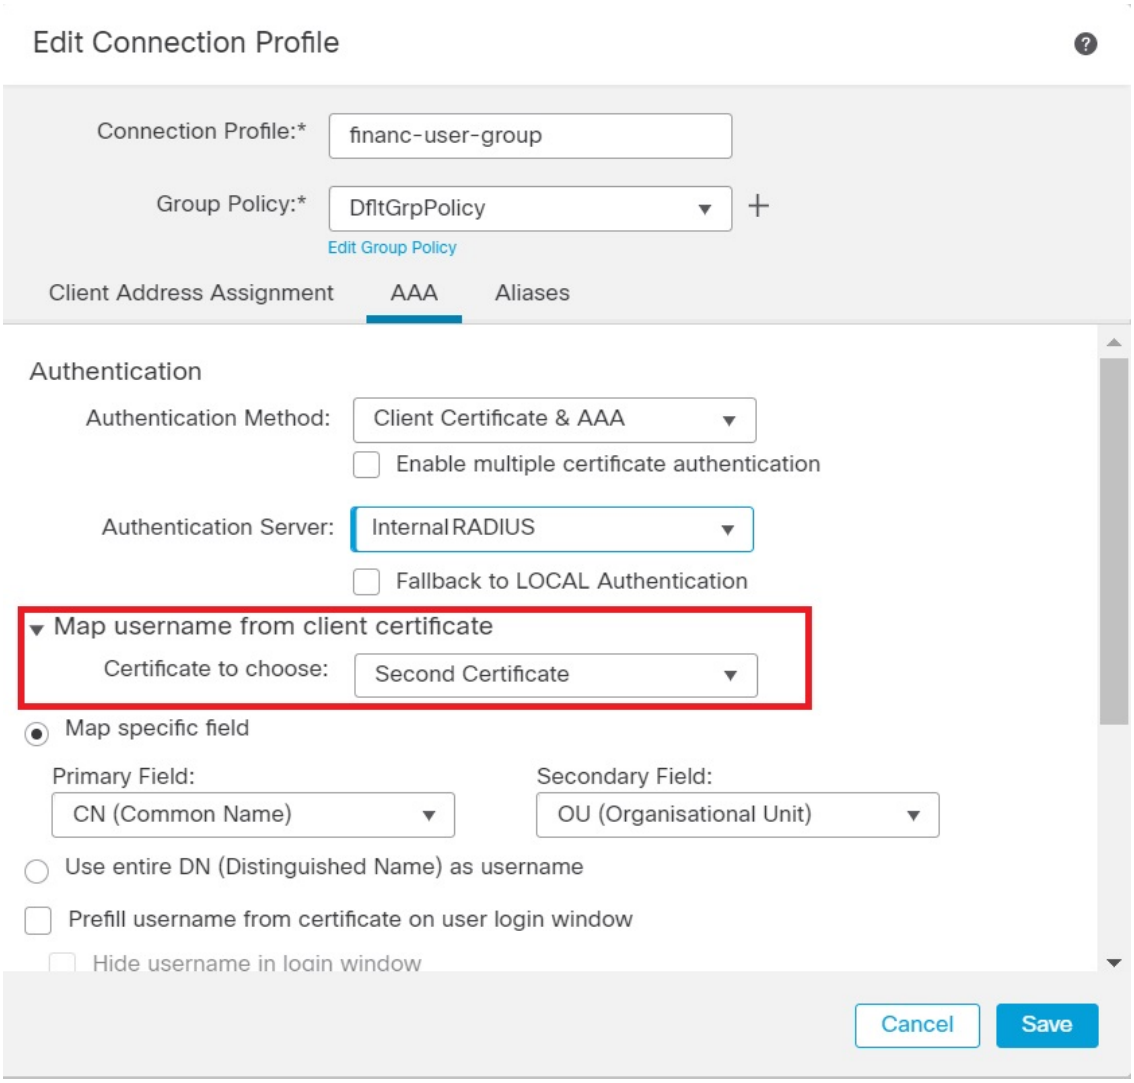

- **Client Certificate Only**—User is authenticated using client certificate. Client certificate must be configured on VPN client endpoints. By default, the username is derived from client certificate fields CN & OU respectively. In case, the username is specified in other fields in the client certificate, use 'Primary' and 'Secondary' field to map appropriate fields.
- **Client Certificate & AAA**—User is authenticated using both the types of authentication, AAA and client certificate.

5. Select **Enable multiple certificate authentication**.
6. Select **Map username from client certificate** and select a certificate from the **Certificate to choose** drop-down to choose the username for the VPN session from the machine certificate or user certificate.
   - **First Certificate** —Map the username from the Machine Certificate.
   - **Second Certificate**—Map the username from the User certificate to authenticate the VPN user.

7. Configure the required connection profile settings and remote access VPN settings.
8. Save the connection profile and remote access VPN policy. Deploy the remote access VPN on Firewall Threat Defense.

For information about remote access VPN AAA settings, see Configure AAA Settings for Remote Access VPN, on page 1597.

Certificate Configuration in DAP

You can also configure certificate criteria attributes in a DAP record. The user and machine certificate received from the VPN client during multiple-certificate authentication is loaded into dynamic access policy (DAP) to allow policies to be configured based on the field of the certificate. You can make policy decisions based on the fields of a certificate used to authenticate that connection attempt.

1. Choose **Devices** > **Dynamic Access Policy**.
2. Edit an existing DAP policy or create a new one and then edit the policy.
3. Choose an existing DAP record, or create a new one and then edit the record.
4. Select **Endpoint Criteria** > **Certificate**.
5. Select the Match Criteria **All** or **Any**.
6. Click **Add** to add certificate attributes.

Figure 438:

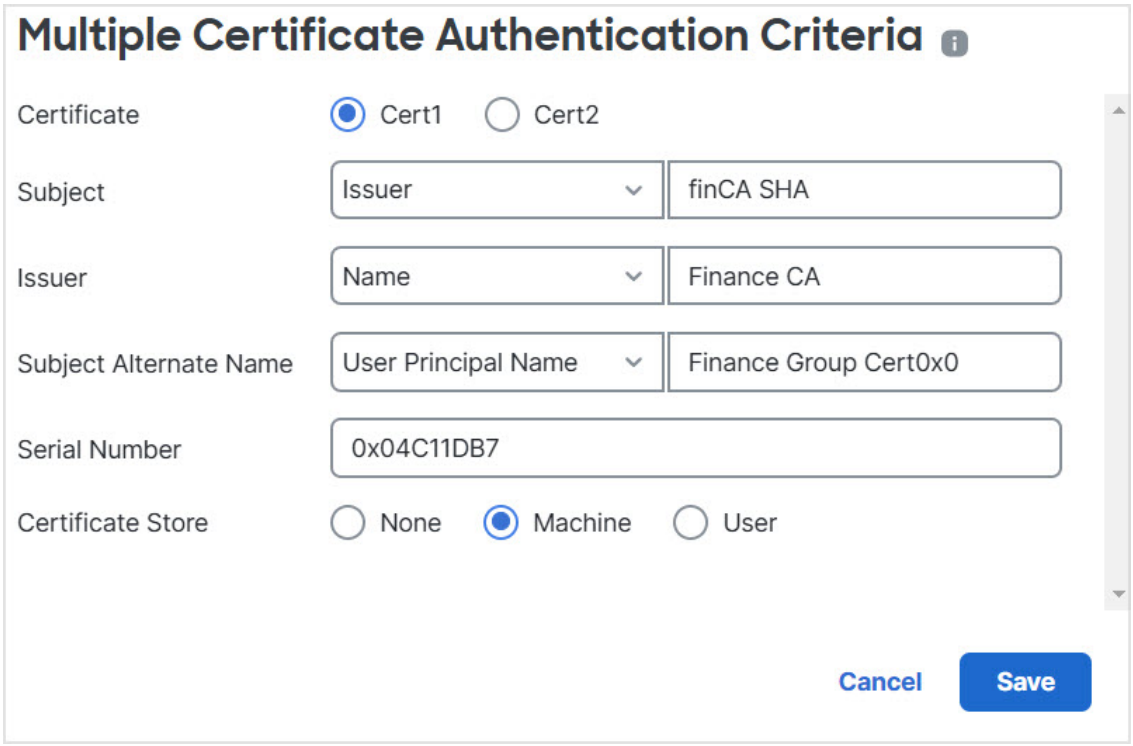

7. Select the certificate, **Cert1** or **Cert2**.

8. Select the **Subject** and specify the certificate subject value.
9. Select the **Issuer** and specify the certificate issuer name.
10. Select the **Subject Alternate Name** and specify the alternate name for the subject.
11. Specify the **Serial Number**.
12. Select the **Certificate Store**: None, Machine, or User.

    This option adds a condition to check for the store from which the certificate is picked on the endpoint.
13. Click **Save** to complete the certificate criteria settings.

    Configure the required DAP record settings and then associate the DAP with the remote access VPN.

For more information about DAP, see Dynamic Access Policies , on page 1689.

History for Remote Access VPNs

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Geolocation-based RA VPN	7.7	7.7	You can now allow or block remote access VPN connections based on country or region. Connections that don't meet your location-based criteria are blocked before authentication and logged for auditing purposes. New/Modified Screens: • **Objects** > **Object Management** > **Access List** > **Service Access** • **Devices** > **VPN** > **Remote Access**
IPsec flow offload on the VTI loopback interface	7.4	Any	On the Secure Firewall 3100 and Secure Firewall 4200 devices, IPsec flow offload is enabled automatically on the VTI loopback interface.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Secure Client Customizations	7.4	Any	You can configure Secure Client customizations and deploy them to the VPN headend. The threat defense distributes these customizations to the endpoint when a user connects from the Secure Client: • GUI text and messages • Icons and images • Scripts • Binaries • Customized Installer Transforms • Localized Installer Transforms New/Modified Screens: • **Objects** > **Object Management** > **VPN** > **Secure Client Customization** • **Devices** > **Remote Access** > **Advanced** > **Secure Client Customization**
WAN Summary Dashboard	7.4	Any	The WAN Summary dashboard provides a snapshot of your WAN devices and their interfaces. New/Modified Screen: **Overview** > **Dashboard** > **WAN Summary**
SAML with Certificate Support	7.2	Any	We have updated the remote access VPN configuration wizard to support user authentication with Certificate and SAML. You can configure a remote access VPN to authenticate machine or user certificate before a SAML authentication is initiated.
IPsec flow offload.	7.2	Any	On the Secure Firewall 3100, IPsec flows are offloaded by default. After the initial setup of an IPsec site-to-site VPN or remote access VPN security association (SA), IPsec connections are offloaded to the field-programmable gate array (FPGA) in the device, which should improve device performance. You can change the configuration using FlexConfig and the **flow-offload-ipsec** command.
Multiple IDP trustpoint support	7.1	Any	Secure Firewall Management Center supports multiple identity provider trustpoints with Microsoft Azure that can have multiple applications for the same Entity ID, but a unique identity certificate.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
AnyConnect VPN SAML External Browser	7.1	Any	You can now configure AnyConnect VPN SAML External Browser to enable additional authentication choices, such as password less authentication, WebAuthN, FIDO, SSO, U2F, and an improved SAML experience due to the persistence of cookies. When you use SAML as the primary authentication method for a remote access VPN connection profile, you can elect to have the AnyConnect client use the client's local browser instead of the AnyConnect embedded browser to perform the web authentication. This option enables single sign-on (SSO) between your VPN authentication and other corporate logins. Also choose this option if you want to support web authentication methods, such as biometric authentication and Yubikeys, that cannot be performed in the embedded browser. We updated the remote access VPN connection profile wizard to allow you to configure the **SAML Login Experience**.
Multi-Certificate Authentication	7.0	Any	Secure Firewall Management Center now supports multiple certificate-based authentication for Firewall Threat Defense to validate the machine or device certificate, to ensure that the device is a corporate-issued device in addition to authenticating the user's identity certificate to allow VPN access using AnyConnect client.
VPN Load balancing	7.0	Any	VPN load balancing logically group two or more devices and distributes remote access VPN sessions among the grouped devices equally without considering throughput and other traffic parameters.
AnyConnect Custom Attributes	7.0	Any	Secure Firewall Management Center now supported AnyConnect custom attributes and provides an infrastructure to configure the AnyConnect client feature without adding hard-coded support for these features on Firewall Threat Defense.
Local User Authentication	7.0	Any	You can now configure and manage users locally on Firewall Threat Defense using the Secure Firewall Management Center web interface, and configure the local users for primary and secondary remote access VPN authentication.
Selective Policy Deployment	7.0	Any	You can now choose to include or exclude changes to remote access VPN and site-to-site VPN configurations during the deployment.
Support for AnyConnect Modules Configuration	6.7	Any	Secure Firewall Management Center now supports configuring AnyConnect modules and profiles for additional security.
Support for LDAP Authorization	6.7	Any	You can configure LDAP authorization for remote access VPN using the Secure Firewall Management Center.
SAML single sign-on support for remote access VPN	6.7	Any	You can configure a SAML 2.0 server as the single sign-on authentication server for remote access VPNs.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
AnyConnect Management VPN tunnel support	6.7	Any	Firewall Threat Defense remote access VPN supports configuring AnyConnect Management VPN tunnel that allows VPN connectivity to endpoints when the corporate endpoints are powered on, without the VPN users connecting to the VPN.
Support for Datagram Transport Layer Security (DTLS) 1.2	6.6	Any	DTLS 1.2 is now part of the default SSL cipher group and it can be configured along with TLS 1.2.

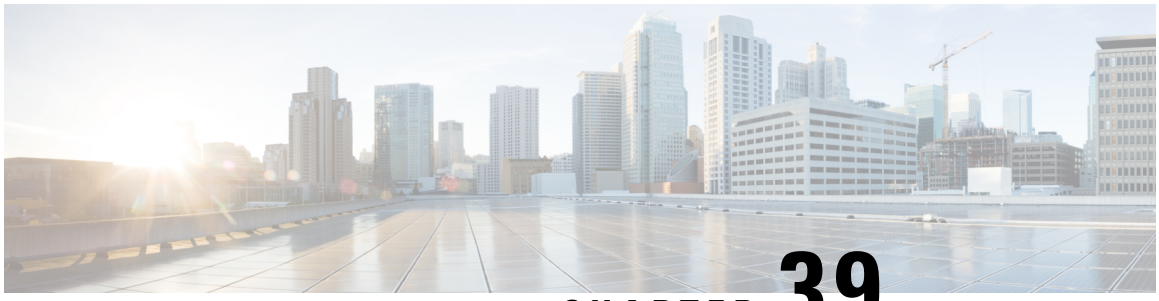

CHAPTER 39

Dynamic Access Policies

Dynamic access policies (DAP) enable you to configure authorization that addresses the dynamics of VPN environments. You create a dynamic access policy by setting a collection of access control attributes that you associate with a specific user tunnel or session. These attributes address issues of multiple group membership and endpoint security.

- About Secure Firewall Threat Defense Dynamic Access Policy, on page 1689
- Prerequisites for Dynamic Access Policy , on page 1691
- Guidelines and Limitations for Dynamic Access Policies, on page 1691
- Configure a Dynamic Access Policy (DAP), on page 1691
- Associate Dynamic Access Policy with Remote Access VPN, on page 1701
- History for Dynamic Access Policy, on page 1702

About Secure Firewall Threat Defense Dynamic Access Policy

VPN gateways operate in dynamic environments. Multiple variables can affect each VPN connection. For example, intranet configurations that frequently change, the various roles each user inhabits within an organization, and log in attempts from remote access sites with different configurations and levels of security. The task of authorizing users is much more complicated in a VPN environment than it is in a network with a static configuration.

You can create a dynamic access policy by setting a collection of access control attributes that you associate with a specific user tunnel or session. These attributes address issues of multiple group memberships and endpoint security. The Firewall Threat Defense grants access to a particular user for a particular session according to the policies you define. The Firewall Threat Defense device generates a DAP during user authentication by selecting or aggregating attributes from one or more DAP records. The device then selects these DAP records based on the endpoint security information of the remote device and AAA authorization information for the authenticated user. Then the device applies the DAP record to the user tunnel or session.

Hierarchy of Policy Enforcement of Permissions and Attributes in Firewall Threat Defense

The Firewall Threat Defense device supports applying user authorization attributes, also called user entitlements or permissions, to VPN connections. The attributes are applied from a DAP on the Firewall Threat Defense, external authentication server and/or authorization AAA server (RADIUS) or from a group policy on the Firewall Threat Defense device.

If the Firewall Threat Defense device receives attributes from all sources, the device evaluates, merges, and applies the attributes to the user policy. If there are conflicts between attributes coming from the DAP, the AAA server, or the group policy, the attributes from the DAP always take precedence.

The Firewall Threat Defense device applies attributes in the following order:

Figure 439: Policy Enforcement Flow

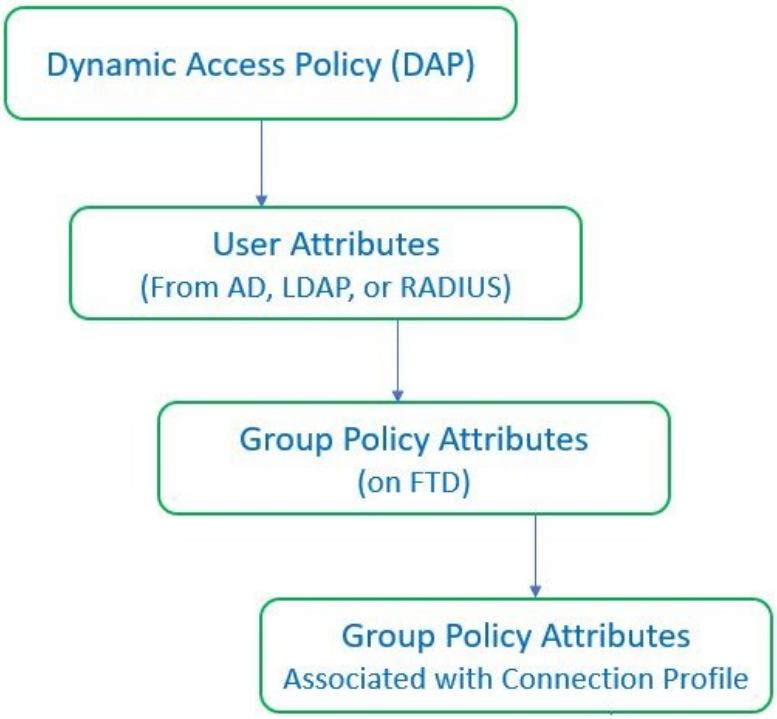

1. **DAP attributes on the FTD**—The DAP attributes take precedence over all others.
2. **User attributes on the external AAA server**—The server returns these attributes after successful user authentication and/or authorization.
3. **Group policy configured on the FTD** —If a RADIUS server returns the value of the RADIUS Class attribute IETF-Class-25 (OU= group-policy) for the user, the Firewall Threat Defense device places the user in the group policy of the same name and enforces any attributes in the group policy that are not returned by the server.
4. **Group policy assigned by the Connection Profile (also known as Tunnel Group)**—The Connection Profile has the preliminary settings for the connection, and includes a default group policy that is applied to the user before authentication.

Note The Firewall Threat Defense device does not support inheriting system default attributes from the default group policy, *DfltGrpPolicy*. For the user session, the device uses the attributes on the group policy that you assign to the connection profile, unless the user attributes or the group policy from the AAA server overrides them.

Prerequisites for Dynamic Access Policy

Licensing:

- Firewall Threat Defense must have at least one of the following Secure Client licenses:
  - Secure Client Premier
  - Secure Client Advantage
  - Secure Client VPN Only
- The Firewall Threat Defense Essentials license must allow export-controlled functionality.

Guidelines and Limitations for Dynamic Access Policies

- Matching of AAA attributes in a DAP will work only if a AAA server is configured to return the correct attributes when authenticating or authorizing a remote access VPN session.
- Minimum Secure Client and Secure Firewall Posture package version supported for DAP is 4.6. But it is highly recommended to use the latest version of Secure Client.
- DAP does not support clustering or multi-instance mode.
- DAP condition with an assigned IPv4 or IPv6 address does not work for local authentication.

Configure a Dynamic Access Policy (DAP)

Create a Dynamic Access Policy

Before you begin

Ensure that you have the Secure Firewall Posture package before you configure the dynamic access policy. You can add the Secure Firewall Posture file at **Objects** > **Object Management** > **VPN** > **Secure Client File**.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**, and click **Create Dynamic Access Policy**.

Step 2 Specify the **Name** for the DAP policy and an optional **Description**.

Step 3 Choose the **Secure Firewall Posture Package** from the drop-down list.

Step 4 Click **Save**.

What to do next

To configure DAP record, see Create Dynamic Access Policy Record

Create a Dynamic Access Policy Record

A dynamic access policy (DAP) can contain multiple DAP records, where you configure user and endpoint attributes. You can prioritize the DAP records within a DAP so that the Firewall Threat Defense can select and sequence the required criteria when a user attempts VPN connection.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.

Step 2 Edit an existing dynamic access policy or click **Create Dynamic Access Policy** to create a new one and then edit the policy.

Step 3 Click **Create DAP Record**.

Step 4 Click the **General** tab.

Step 5 Specify the **Name** for the DAP record.

Step 6 Enter the **Priority** for the DAP record.

The lower the number, the higher the priority.

Step 7 Select one of the following actions to take when a DAP record matches:

- **Continue**—Applies access policy attributes to the session. If any, the next DAP record (next policy line with less priority) is then evaluated.
- **Terminate**—Terminates the session.
- **Quarantine**—Quarantines the connection.

Step 8 Check the **Display User Message on Criterion Match** check-box and add the user message.

The Firewall Threat Defense displays this message to the user when the DAP record matches.

Step 9 Check the **Apply a Network ACL on Traffic** check box and select the access control list from the drop-down.

Step 10 Check the **Apply one or more Secure Client Custom Attributes** check box and select the custom attributes object from the drop-down.

Step 11 Click **Save**.

Configure Posture Assessment Criteria

For a DAP policy, you can configure file, process, or registry endpoint attributes with unique endpoint IDs. These IDs can be used as endpoint criteria in Lua scripts to configure a DAP record.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.

Step 2 Click **Create Dynamic Access Policy** to create a new DAP policy. and then edit the policy.

Step 3 Click the edit icon adjacent to the DAP policy.

Step 4 Click **Add Posture Assessment Criteria**.

Step 5 Do one of the following:

- Configure a file endpoint attribute:
  - **a.** Click the **File** radio button.
  - **b.** In the **Endpoint ID** field, enter a unique ID for the file. It can be a string or a number.
  - **c.** In the **File Path** field, specify the file path.
- Configure a registry endpoint attribute:
  - **a.** Click the **Registry** radio button.
  - **b.** In the **Endpoint ID** field, enter a unique ID for the registry. It can be a string or a number.
  - **c.** In the **Entry Path** field, specify the file path.
- Configure a process endpoint attribute:
  - **a.** Click the **Process** radio button.
  - **b.** In the **Endpoint ID** field, enter a unique ID for the process. It can be a string or a number.
  - **c.** In the **Process Name** field, specify the process name.

Note
You cannot edit an Endpoint ID once it is saved.

Step 6 Click **Save**.

What to do next

You can use the endpoint IDs to configure advanced posture assessment criteria using Lua scripts. For more information, see Configure Advanced Settings for DAP, on page 1700.

Configure AAA Criteria Settings for DAP

DAP complements AAA services by providing a limited set of authorization attributes that can override the attributes that AAA provides. The Firewall Threat Defense select DAP records based on the AAA authorization information for the user and posture assessment information for the session. The Firewall Threat Defense can choose multiple DAP records depending on this information, which it then aggregates to create DAP authorization attributes.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.

Step 2 Edit an existing DAP policy or create a new one and then edit the policy.

Step 3 Select a DAP record or create a new one, and edit the DAP record.

Step 4 Click **AAA Criteria**.

Step 5 Select one of the **Match criteria between sections**.

- Any—Matches any of the criteria.
- All—Matches all the criteria.
- None—Matches none of the set criteria.

Step 6 Click **Add** to add the required **Cisco VPN Criteria**.

Cisco VPN criteria include attributes for group policy, assigned IPv4 address, assigned IPv6 address, connection profile, username, username 2, and SCEP required.

a) Select an attribute and specify the **Value**.

b) Click **Add another criteria** to add more criteria.

c) Click **Save**.

SCEP Required

Step 7 Select **LDAP Criteria**, **RADIUS Criteria**, or **SAML Criteria** and specify the **Attribute ID** and **Value**.

Step 8 Click **Save**.

Configure Endpoint Attribute Selection Criteria in DAP

Endpoint attributes contain information about the endpoint system environment, posture assessment results, and applications. The Firewall Threat Defense dynamically generates a collection of endpoint attributes during session establishment and stores these attributes in a database that is associated with the session. Each DAP record specifies the endpoint selection attributes that must be satisfied for the Firewall Threat Defense to choose it for a session. The Firewall Threat Defense selects only DAP records that satisfy every condition configured.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**, and click **Create Dynamic Access Policy**.

Step 2 Edit a DAP policy and then DAP record.

Note
Create a DAP policy and DAP record if not done already.

Step 3 Click **Endpoint Criteria** and configure the following endpoint criteria attributes:

Note
You can create multiple instances of each type of endpoint attribute. There is no limit for the number of endpoint attributes for each DAP record.

- Add an Anti-Malware Endpoint Attribute to a DAP
- Add a Device Endpoint Attribute to a DAP
- Add Secure Client Endpoint Attributes to a DAP, on page 1696
- Add a NAC Endpoint Attribute to a DAP
- Add an Application Attribute to a DAP
- Add a Personal Firewall Endpoint Attribute to a DAP
- Add an Operating System Endpoint Attribute to a DAP
- Add a Process Endpoint Attribute to a DAP
- Add a Registry Endpoint Attribute to a DAP
- Add a File Endpoint Attribute to a DAP
- Add Multiple Certificate Authentication Attributes to DAP

Step 4 Click **Save**.

Add an Anti-Malware Endpoint Attribute to a DAP

Procedure

Step 1 Edit a DAP record and select **Endpoint Criteria** > **Anti-Malware**.

Step 2 Select the Match Criteria **All** or **Any**.

Step 3 Click **Add** to add anti-malware attributes.

Step 4 Click **Installed** to indicate whether the selected endpoint attribute and its accompanying qualifiers are installed or not installed.

Step 5 Choose **Enabled** or **Disabled** to activate or deactivate real-time malware scanning.

Step 6 Select the name of the anti-malware **Vendor** from the list.

Step 7 Select the anti-malware **Product Description**.

Step 8 Choose the **Version** of the anti-malware product.

Step 9 Specify the number of days since the **Last Update**.

You can indicate that an anti-malware update must occur in less than (<) or more than (>) the number of days you specify.

Step 10 Click **Save**.

Add a Device Endpoint Attribute to a DAP

Procedure

Step 1 Edit a DAP record and choose **Endpoint Criteria** > **Device**.

Step 2 Select the Match Criteria **All** or **Any**.

Step 3 Click **Add** and select the = or ≠ operator to check the attribute to be equal to or not equal to the value you enter for the following attributes:

- **Host Name**—Hostname of the device you are testing for. Use the computer's host name only, not the fully qualified domain name (FQDN).
- **MAC Address**—MAC address of the network interface card you are testing for. The address must be in the format xxxx.xxxx.xxxx where x is a hexadecimal character.
- **BIOS Serial Number**—BIOS serial number value of the device you are testing for. The number format is manufacturer-specific.
- **Port Number**—Listening port number of the device.
- **Secure Desktop Version**—Version of the Host Scan image running on the endpoint.
- **OPSWAT Version**—The OPSWAT client version.
- **Privacy Protection**—None, Cache cleaner, Secure Desktop.
- **TCP/UDP Port Number**—TCP or UDP port in the listening state that you are testing for.

Step 4 Click **Save**.

Add Secure Client Endpoint Attributes to a DAP

Procedure

Step 1 Edit a DAP record and select **Endpoint Criteria** > **Secure Client**.

Step 2 Select the Match Criteria **All** or **Any**.

Step 3 Click **Add** and select the = or ≠ operator to check the attribute to be equal to or not equal to the value you enter.

Step 4 Select the **Client Version** and **Platform**.

Step 5 Select the **Platform Version**, and specify the **Device Type** and **Device Unique ID**.

Step 6 Add the **MAC Addresses** the MAC Address Pool.

Note
The MAC Address must be in the format XX-XX-XX-XX-XX-XX, where each X is a hexadecimal character. You can click **Add another MAC Address** to add more addresses.

Step 7 Click **Save**.

Add NAC Endpoint Attributes to a DAP

Procedure

Step 1 Edit a DAP record and select **Endpoint Criteria** > **NAC**.

Step 2 Select the Match Criteria **All** or **Any**.

Step 3 Click **Add** to add NAC attributes.

Step 4 Set the operator to be equal to = or not equal to ≠ the posture token string. Enter the posture token string in the **Posture Status** box.

Step 5 Click **Save**.

Add an Application Attribute to a DAP

Procedure

Step 1 Edit a DAP record and select **Endpoint Criteria** > **Application**.

Step 2 Select the Match Criteria **All** or **Any**.

Step 3 Click **Add** to add application attributes.

Step 4 Choose equals (=) or does not equal (≠) and specify the **Client Type** to indicate the type of remote access connection.

Step 5 Click **Save**.

Add a Personal Firewall Endpoint Attribute to a DAP

Procedure

Step 1 Edit a DAP record and select **Endpoint Criteria** > **Personal Firewall**.

Step 2 Select the Match Criteria **All** or **Any**.

Step 3 Click **Add** to add personal firewall attributes.

Step 4 Click **Installed** to indicate whether the personal firewall endpoint attribute and its accompanying qualifiers (fields below the Name/Operation/Value column) are installed or not installed.

Step 5 Choose **Enabled** or **Disabled** to activate or deactivate firewall protection.

Step 6 Select the name of the firewall **Vendor** from the list.

Step 7 Select the firewall **Product Description**.

Step 8 Select the equals (=) or does not equal (≠) operator and choose the **Version** of the personal firewall product.

Step 9 Click **Save**.

Add an Operating System Endpoint Attribute to a DAP

Procedure

Step 1 Edit a DAP record and select **Endpoint Criteria** > **Operating System** .
Step 2 Select the Match Criteria **All** or **Any**.
Step 3 Click **Add** to add endpoint attributes.
Step 4 Select the equals (=) or does not equal (≠) operator and then select the **Operating System**.
Step 5 Select the equals (=) or does not equal (≠) operator and then specify the operating system **Version**.
Step 6 Click **Save**.

Add a Process Endpoint Attribute to a DAP

Procedure

Step 1 Edit a DAP record.
Step 2 Click the **Endpoint Criteria** tab.
Step 3 Click **Process**.
Step 4 Select the **Match Criteria** as **All** or **Any**.
Step 5 Click + to add the process attributes.
Step 6 Select **Exists** or **Does not exist**.
Step 7 Specify the **Process Name**.
Step 8 From the **Endpoint ID** drop-down list, choose the ID for the process or click + to configure a posture assessment criteria for the process. For more information, see Configure Posture Assessment Criteria, on page 1692.
Step 9 Click **Exists** or **Does not exist**.
Step 10 Click **Save**.

Add a Registry Endpoint Attribute to a DAP

Scanning for registry endpoint attributes applies to Windows operating systems only.

Before you begin

Before configuring a Registry endpoint attribute, define the registry key for which you want to scan in the Host Scan window for Cisco Secure Desktop.

Procedure

Step 1 Edit a DAP record.
Step 2 Click the **Endpoint Criteria** tab.
Step 3 Click **Registry**.
Step 4 Select the **Match Criteria** as **All** or **Any**.
Step 5 Click + to add registry attributes.
Step 6 Select the **Entry Path** for the registry and specify the path.
Step 7 From the **Endpoint ID** drop-down list, choose the ID for the registry or click + to configure a posture assessment criteria for the registry. For more information, see Configure Posture Assessment Criteria, on page 1692.
Step 8 Choose the existence of the registry, **Exists** or **Does not exist**.
Step 9 Select the registry **Type** from the list.
Step 10 Select the equals (=) or does not equal (≠) operator and enter the **Value** of the registry key.
Step 11 Select **Case insensitive** to disregard the case of the registry entry while scanning.
Step 12 Click **Save**.

Add a File Endpoint Attribute to a DAP

Procedure

Step 1 Edit a DAP record.
Step 2 Click the **Endpoint Criteria** tab.
Step 3 Click **File**.
Step 4 Select the Match Criteria **All** or **Any**.
Step 5 Click + to add file attributes.
Step 6 Specify the **File Path**.
Step 7 From the **Endpoint ID** drop-down list, choose the ID for the file or click + to configure a posture assessment criteria for the file. For more information, see Configure Posture Assessment Criteria, on page 1692.
Step 8 Choose **Exists** or **Does not exist** to indicate the presence of the file.
Step 9 Select less than (<) or greater than (>) and specify the **Last Modified** days for the file.
Step 10 Select the equal to (=) or not equal to ≠ operator and enter the **Checksum**.
Step 11 Click **Save**.

Add Certificate Authentication Attributes to a DAP

You can index each certificate to allow referencing to any of the received certificates, by the configured rules. Based on these certificate fields, you can configure DAP rules to allow or disallow connection attempts.

Procedure

Step 1 Edit a DAP record and select **Endpoint Criteria** > **Certificate**.

Step 2 Select the Match Criteria **All** or **Any**.

Step 3 Click **Add** to add certificate attributes.

Step 4 Select the certificate **Cert1** or **Cert2**.

Step 5 Select the **Subject** and specify the subject value.

Step 6 Select the **Issuer** and specify the issuer value.

Step 7 Select the **Subject Alternate Name** and specify the subject value.

Step 8 Specify the **Serial Number**.

Step 9 Choose the **Certificate Store**: None, Machine, or User.

The VPN client sends the certificate store information.

Step 10 Click **Save**.

Configure Advanced Settings for DAP

You can use the **Advanced** tab for adding selection criteria other than what is possible in the AAA and endpoint attribute areas. For example, while you can configure the Firewall Threat Defense to use AAA attributes that satisfy any, all, or none of the specified criteria, the endpoint attributes are cumulative, and must satisfy all. To let the security appliance employ one endpoint attribute or another, you must create appropriate logical expressions in Lua and enter them here.

Procedure

Step 1 Choose **Devices** > **VPN** > **Dynamic Access Policy**.

Step 2 Edit a DAP policy and then edit a DAP record.

Note
Create a DAP policy and DAP record if not done already.

Step 3 Click the **Advanced** tab.

Step 4 Select **AND** or **OR** as the match criteria to use in the DAP configuration.

Step 5 Add the Lua script in the **Lua script for advanced attribute matching** field.

Step 6 To use the endpoint criteria ID in your Lua script:

a. Place the cursor at the point where you want to insert the endpoint criteria ID.

b. From the **Endpoint Criteria** drop-down list, choose the criteria a

c. Choose the corresponding ID from the adjacent drop-down list.

Example:

In the following example, DAPTESTFILE, LIBAGENT, vpnagent, and DUOAGENT were inserted in the Lua script:

```
EVAL(endpoint.file["DAPTESTFILE"].exists,"EQ","true") or
EVAL(endpoint.file["LIBAGENT"].exists,"EQ","true") and
EVAL(endpoint.process[""vpnagent""].exists,"EQ","true") and
EVAL(endpoint.registry[""DUOAGENT""].exists,"EQ","true")
```

Step 7 Click **Save**.

Associate Dynamic Access Policy with Remote Access VPN

You can associate Dynamic Access Policy (DAP) with remote access VPN policy for the dynamic access policy attributes to match during VPN session authentication and authorization. You can then deploy the remote access VPN on the Firewall Threat Defense.

Procedure

Step 1 Choose **Devices** > **VPN** > **Remote Access**.

Step 2 Click **Edit** next to the remote access VPN policy to which you want to associate dynamic access policy.

Step 3 Click the link in remote access VPN to select the dynamic access policy.

Step 4 Select the policy from the **Dynamic Access Policy** drop-down or click **Create a new Dynamic Access Policy** to configure a new dynamic access policy.

Step 5 Click **OK**.

Step 6 Click **Save** to save the remote access VPN policy.

When the remote access VPN user tries to connect, the VPN checks the configured dynamic access policy records and attributes. VPN creates a dynamic access policy based on the matching dynamic access policy records and takes appropriate action on the VPN session.

History for Dynamic Access Policy

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Easily configure posture assessment criteria for dynamic access policies	7.7	Any	For a DAP policy, you can configure file, process, or registry endpoint attributes with unique endpoint IDs. These IDs can be used as endpoint criteria in Lua scripts to configure a DAP record. New/Modified Screens: • **Devices** > **Dynamic Access Policy** > **Add/Edit Policy** > **Add Posture Assessment Criteria** • **Devices** > **Dynamic Access Policy** > **Add/Edit Policy** > **Add/Edit DAP Record** > **Advanced** > **Endpoint Criteria**
Dynamic Access Policy	7.0	Any	The feature was introduced.

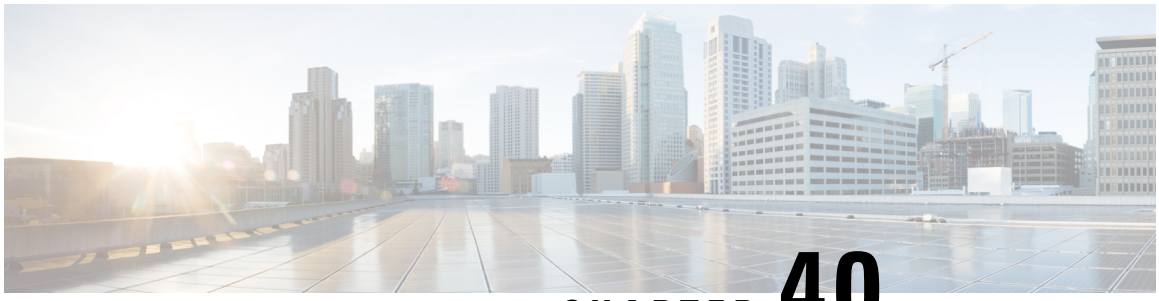

CHAPTER 40

VPN Monitoring and Troubleshooting

This chapter describes Firewall Threat Defense VPN monitoring tools, parameters, and statistics information as well as troubleshooting.

- Site-to-Site VPN Summary Page, on page 1703
- Remote Access VPN Dashboard, on page 1703
- SD-WAN Summary Dashboard, on page 1705
- VPN Session and User Information, on page 1710
- Site to Site VPN Connection Event Monitoring, on page 1710
- VPN Troubleshooting, on page 1712

Site-to-Site VPN Summary Page

You can use the Site-to-Site VPN Summary page to see consolidated information about VPN users, including the current status of users, device types, client applications, user geolocation information, and duration of connections. You can view details of the configured VPN topologies such as VPN interfaces, tunnel status, and so on.

For all VPN topologies, you can edit or delete the topology using the edit and delete buttons. For SASE topology VPNs, you have options to deploy, edit and delete any topology.

Remote Access VPN Dashboard

Remote Access Virtual Private Network (RA VPN) allows remote users to securely connect to your network. The RA VPN dashboard allows you to monitor real-time data from active RA VPN sessions on the devices. You can quickly determine problems related to user sessions and mitigate the problems for your network and users.

RA VPN dashboard (**Overview** > **Dashboards** > **Remote Access VPN**) provides a snapshot of the active RA VPN sessions on the threat defense devices managed by the Firewall Management Center.

The dashboard has the following widgets:

- Active Sessions (Tabular View)
- Active Sessions (Map view)
- Sessions

- Device Identity Certificates

Active Sessions (Tabular View)

This widget provides a tabular view of the active RA VPN users connected. You can view details of the active RA VPN sessions such as username, assigned IP, public IP, login time, VPN gateway (threat defense device), client application, client operating system, connection profile, and group policy. You can use the filter to narrow down your search based on the different criteria. You can also perform the following actions on the individual sessions:

- Terminate a session of a specific user.
- Terminate all sessions of a specific user connected to a specific VPN gateway.
- Terminate all sessions that are connected to a specific VPN gateway.

If a client device supports dual address stack and the RA VPN configuration on the Firewall Threat Defense device allows IPv4 and IPv6 address pools, when a client establishes an RA VPN session with the headend device, it assigns an IPv4 and an IPv6 address to the client's tunnel interface. The RA VPN session has two IP addresses, an IPv4 and an IPv6 address on the threat defense device. The Firewall Management Center shows two sessions for the same user, one with an IPv4 address and another with an IPv6 address, and the session count is two.

Hence, even when there is only a single RA VPN session from a user as per **show vpn-sessiondb l2l filter ipaddress** command on the device, the Firewall Management Center shows two different sessions.

Active Sessions (Map View)

This widget shows an interactive heat map to visualize the location of the users connected through RA VPN sessions on the devices.

- Countries that have user sessions appear in shades of blue.
- Legend of the map provides a scale that indicates the correlation between the number of sessions in a country and the shade of blue for the country.
- Hover the mouse pointer over the map to view the country name and the total number of active user sessions.
- Zoom in, zoom out, and reset options are available.

Sessions

This widget allows you to monitor real-time data from active RA VPN sessions on the devices. You can filter and view the distribution of active RA VPN sessions according to:

- Device: Displays the number of sessions per device.
- Encryption Type: Displays the number of Secure Client SSL or IPsec sessions.
- Secure Client Version: Displays the sessions per Secure Client version.
- Operating System: Displays the sessions per operating systems. For example, Windows, Linux, Mac, Mobile OS, and so on.
- Connection Profile: Displays the sessions per connection profile.

Device Identity Certificates

This widget provides information about the identity certificate expiry of the RA VPN gateways. You can view expired certificates and certificates that are due for expiry within a month. Click **View Details** to view the certificates in the **Device > Certificates** page.

SD-WAN Summary Dashboard

The SD-WAN Summary dashboard (**Overview** > **Dashboards** > **SD-WAN Summary**) provides a snapshot of your WAN devices and their interfaces. This dashboard helps you to:

- Identify issues with the underlay and overlay (VPN) topologies.
- Troubleshoot VPN issues using the existing **Health Monitoring**, **Device Management**, and **Site-to-Site Monitoring** pages.
- Monitor application performance metrics of WAN interfaces. The threat defense steers application traffic based on these metrics.

A WAN device must meet one of the following criteria:

- The device must be a VPN peer.
- The device must have WAN interface.

A WAN interface must meet one of the following criteria:

- The interface has IP address-based path monitoring enabled on it.
- The interface has a Policy Based Routing (PBR) policy with at least one application configured to monitor it.

For more information about PBR policy and path monitoring, see Policy Based Routing, on page 1323.

Click **Uplink Decisions** to view the **VPN Troubleshooting** page. You can view syslogs with ID: 880001. These syslogs show the threat defense interfaces through which it steers traffic based on the configured PBR policy.

To view the above syslogs and to view the data on this dashboard, ensure that you review Prerequisites for Using SD-WAN Summary Dashboard, on page 1705.

For clusters, this dashboard displays application performance metrics of only the control node and not the data nodes.

Prerequisites for Using SD-WAN Summary Dashboard

- You must be an Admin, Security Analyst, or Maintenance user to view this dashboard.
- Threat defense devices must be Version 7.2 or later.
- Enable IP-based path monitoring and HTTP-based application monitoring on the WAN interfaces.
  1. Choose **Devices** > **Device Management**.
  2. Click the edit icon adjacent to the device that you want to edit.

3. Click the edit icon adjacent to the interface that you want to edit.
4. Click the **Path Monitoring** tab.
5. Check the **Enable IP based Monitoring** check box.
6. Check the **Enable HTTP based Application Monitoring** check box.
7. Click **OK**.

- Configure a PBR policy with at least one application configured to monitor it:

1. Choose **Devices** > **Device Management**.
2. Click the edit icon adjacent to the device that you want to edit.
3. Click **Routing**.
4. In the left pane, click **Policy Based Routing**.
5. Click **Add**.
6. From the **Ingress Interface** drop-down list, choose an interface.
7. Click **Add** to configure a forwarding action.
8. Configure the parameters.
9. Click **Save**.

- To view the application performance metrics for the WAN interfaces, you must:
  - Threat defense devices must be Version 7.4.1.
  - Enable data collection from the SD-WAN module in the health policy.
    1. Choose **System** (⚙) > **Health** > **Policy**.
    2. Click the **Edit health policy** icon.
    3. In the **Health Modules** tab, under **SD-WAN**, click the **SD-WAN Monitoring** toggle button.
  - Configure applications for the PBR policies.
    1. Choose **Objects** > **Object Management** > **Access List** > **Extended**.
    2. Click the edit icon adjacent to the access list and add the applications for the PBR policy.
  - Configure the forwarding action for the policy with one of the four application metrics.
    1. Choose **Devices** > **Device Management**.
    2. Click the edit icon adjacent to the device that you want to edit.
    3. Click **Routing**.
    4. In the left pane, click **Policy Based Routing**.
    5. Click the edit icon adjacent to the policy that you want to edit.

6. In the **Edit Policy Based Route** dialog box, click the edit icon adjacent to the corresponding ACL.
7. In the **Edit Forwarding Actions** dialog box, from the **Interface Ordering** drop-down list, choose one of the following options:
   - **Minimal Jitter**
   - **Maximum Mean Opinion Score**
   - **Minimal Round-Trip Time**
   - **Minimal Packet Loss**

   If you choose **Interface Priority** or **Order**, application monitoring is not enabled on the interface.

- Configure ECMP on the WAN interfaces:
  1. Choose **Devices** > **Device Management**.
  2. Click the edit icon adjacent to the device that you want to edit.
  3. Click **Routing**.
  4. In the left pane, click **ECMP**.
  5. Click **Add** and specify a name for the ECMP zone.
  6. Click **Add** to move interfaces from **Available Interfaces** to **Selected Interfaces**.
  7. Click **OK**.
- Ensure that traffic passes through the interface.
- Enable DNS inspection on each WAN device so that the threat defense device can do DNS snooping, and configure the trusted DNS servers:
  1. Choose **Devices** > **Platform Settings**.
  2. Click the edit icon adjacent to the threat defense policy that you want to edit.
  3. In the left pane, click **DNS**.
  4. Click the **DNS Settings** tab.
  5. Check the **Enable DNS name resolution by device** check box.
  6. Click the **Trusted DNS Servers** tab.
  7. Do one of the following:
     - Click the **Trust Any DNS server** toggle button.
     - Under **Specify DNS Servers**, click **Edit** to add trusted DNS servers.

- To view syslogs when you click **Uplink Decisions**, you must:
  - Choose **Devices** > **Platform Settings** and create or edit a threat defense policy.
  - In the left pane, click **Syslog**.

- Click the **Logging Setup** tab.
- Check the **Enable Logging** check box to turn on the data plane system logging for the threat defense device.
- Click the **All Logs** radio button to enable logging of all the troubleshooting syslog messages.

  or

  Click the **VPN Logs** radio button to enable logging of only the VPN troubleshooting messages.
- Click **Save**.

Monitor WAN Devices and Interfaces Using the SD-WAN Summary Dashboard

The SD-WAN Summary dashboard has the following widgets under the **Overview** tab:

- Top Applications, on page 1708
- WAN Connectivity, on page 1708
- VPN Topology, on page 1709
- WAN Interface Throughput, on page 1709
- Device Inventory, on page 1709
- WAN Device Health, on page 1709

Top Applications

This widget displays the top 10 applications ranked according to throughput.

You can choose a time range for the widget data from the **Show Last** drop-down list. The range is 15 minutes to two weeks.

WAN Connectivity

This widget provides a summary of the WAN interfaces statuses. It shows the number of WAN interfaces that are in the **Online**, **Offline** or **No Data** states. Note that you cannot monitor subinterfaces using this widget.

Click **View All Interfaces** to view more details about the interfaces in the health monitor page.

If a WAN interface is in the **Offline** or **No Data** state, you can troubleshoot it from the health monitor page:

1. In the **Monitoring** pane, expand **Devices**.
2. Click the corresponding WAN device to view the device-specific health details.
3. Click the **Interface** tab to view the interface status and aggregate traffic statistics for a specific time.

   Alternatively, you can click **View System & Troubleshoot Details**. The health monitor page is displayed with all the necessary details.

VPN Topology

This widget provides a summary of the site-to-site VPN tunnel statuses. It shows the number of **Active**, **Inactive**, and **No Active Data** VPN tunnels.

Click **View All Connections** to view the VPN tunnel details in the **Site-to-site VPN Monitoring** dashboard.

If the tunnels are in the **Inactive** or **No Active Data** state, you can troubleshoot using the **Site-to-site VPN Monitoring** dashboard. In the **Tunnel Status** widget, hover your cursor over a topology, click **View** ( ) and do one of the following:

- Click the **CLI Details** tab to view the details of the VPN tunnels.
- Click the **Packet Tracer** tab to use the packet tracer tool for the topology.

WAN Interface Throughput

This widget monitors the average throughput of the WAN interfaces during the chosen time period.

The interface throughput is classified into four bands. These details aid in cost planning and resourcing. You can choose a time range for the widget data from the **Show Last** drop-down list. The range is from 15 minutes to two weeks.

Click **View Health Monitoring** to view more details about the interface in the health monitor page.

Device Inventory

This widget lists all the managed WAN devices and groups them according to the model.

Click **View Device Management** to view more details about the device in the **Device Management** page.

WAN Device Health

This widget displays the device count according to the health of the WAN devices. You can view the number of devices with errors, warnings, or those that are in **Disabled** state.

Click **View Health Monitoring** to view the alarms, and quickly identify, isolate, and resolve issues.

If the health of a device is affected, you can troubleshoot it from the health monitor page.

1. In the **Monitoring** pane, expand **Devices**.
2. Click the corresponding WAN device to view the device-specific health details.
3. Click **View System & Troubleshoot Details**. The health monitor page is displayed with all the necessary details.

A device can be in **Disabled** state for multiple reasons, including the following:

- Management interface is disabled.
- Device is powered off.
- Device is being upgraded.

Monitor Application Performance Metrics of WAN Interfaces Using the SD-WAN Summary Dashboard

Under the **Application Monitoring** tab, you can select a WAN device and view the application performance metrics for the corresponding WAN interfaces. These metrics include Jitter, Round Trip Time (RTT), Mean Opinion Score (MOS), and Packet Loss.

By default, the metrics data is refreshed every 5 minutes. You can change the refresh time; the range is from 5 to 30 minutes. You can view the metrics in tabular and graphical formats. For each WAN interface, the latest metric value appears in the table. For graphical data, you can choose a time interval of up to 24 hours to view the metrics data for the corresponding WAN interfaces.

VPN Session and User Information

The system generates events that communicate the details of user activity on your network, including VPN-related activity. The system monitoring capabilities enable you to determine quickly whether remote access VPN problems exist and where they exist. You can then apply this knowledge and use your network management tools to reduce or eliminate problems for your network and users. Optionally, you can log out remote access VPN users as needed.

Viewing Remote Access VPN Active Sessions

Analysis > **Users** > **Active Sessions**

Lets you view the currently logged-in VPN users at any given point in time with supporting information such as the user name, login duration, authentication type, assigned/public IP address, device details, client version, endpoint information, throughput, bandwidth consumed group policy, tunnel group and so on. The system allows you to filter current user information, log users out, and delete users from the summary list.

Note If you configure your VPN in a high-availability deployment, the device name displayed against active VPN sessions can be the primary or secondary device that identified the user session.

Viewing Remote Access VPN User Activity

Analysis > **Users heading** > **User Activity**

Lets you view the details of user activity on your network. The system logs historical events and includes VPN-related information such as connection profile information, IP address, geolocation information, connection duration, throughput, and device information.

Site to Site VPN Connection Event Monitoring

The site-to-site VPN connection event allows you to know if the VPN encrypts or do not encrypts the connection and helps you to troubleshoot connectivity issues, especially in multi-hop VPN deployments. The event

dashboard of the Firewall Management Center displays the IP address of the VPN peer (peer's IKE address) which encrypts or decrypts the traffic and displays the VPN action as follows:

- If the connection is decrypted by the VPN, the column **Decrypt Peer** displays the IP address of the peer's address from where the flow was received and displays **Decrypt** as the VPN action.
- If the connection is encrypted by the VPN, the column **Encrypt Peer** displays the IP address of the VPN peer to which the flow is sent and displays **Encrypt** as the VPN action.
- If the VPN server cascades the connection, it gets decrypted on one tunnel and gets re-encrypted on another tunnel. In this case, both **Encrypt Peer** and **Decrypt Peer** IP addresses get appears in the event. The column **VPN Action** displays **VPN Routing** as the action to indicate that the connection transit through the VPN server.

If you enable the bypass Access Control Policy for decrypted traffic (sysopt permit-vpn) option, the system bypasses the Access Control Policy and do not log events for decrypted traffic. This option is disabled by default and all decrypted traffic in the VPN tunnel undergoes ACL inspection.

View Site to Site VPN Connection Events

Access the connection event viewer of the Firewall Management Center to know if the VPN encrypts or do not encrypts the connection traffic and to retrieve the VPN peer details.

Before you begin

Ensure that you enable logging of connection events at the beginning of connection and at the end of connection in the access control rule.

Procedure

Step 1 Choose **Analysis** > **Connections**.

Step 2 Go to **Table View of Connection Events** tab.

Step 3 In the table view of events, multiple fields are hidden by default. To change the fields that appear, click the **x** icon in any column name to display a field chooser.

Step 4 Choose the following columns:

- Decrypt Peer
- Encrypt Peer
- VPN Action

Step 5 Click **Apply**.

See *Connection and Security-Related Connection Events* in the Secure Firewall Management Center Administration Guide for more information on the connection events.

VPN Troubleshooting

This section describes VPN troubleshooting tools and debug information.

System Messages

The Message Center is the place to start your troubleshooting. This feature allows you to view messages that are continually generated about system activities and status. To open the Message Center, click **System Status**, located to the immediate right of the **Deploy** button in the main menu.

VPN System Logs

You can enable logging of VPN troubleshoot syslogs for Firewall Threat Defense devices. Logging information can help you identify and isolate network or device configuration problems. When you enable VPN logging, the Firewall Threat Defense devices send VPN syslogs to the Firewall Management Center.

All VPN syslogs appear with a default severity level **errors** or a higher severity (unless changed). You can manage the VPN logging through Firewall Threat Defense platform settings. You can adjust the message severity level by editing the **VPN Logging Settings** in the Firewall Threat Defense platform settings policy for targeted devices. See Configure Syslog Logging for Firewall Threat Defense Devices, on page 990 for details on enabling VPN logging, configuring syslog servers, and viewing the system logs.

From the Troubleshooting Logs table (**Devices** > **Troubleshoot** > **Troubleshooting Logs**), you can view and analyze the VPN syslog messages to identify and isolate issues with your network and device configuration.

We recommend that you set the logging level of the VPN logs as level 3 (Errors). Setting the VPN logging level to level 4 and above (Warnings, Notifications, Informational or Debugging) could overload the Firewall Management Center.

Note When you configure a device with site-to-site or remote access VPN, it automatically enables sending VPN syslogs to the Firewall Management Center by default.

Debug Commands

This section explains how you use debug commands to help you diagnose and resolve VPN-related problems. The commands described here are not exhaustive, this section include commands according to their usefulness in assisting you to diagnose VPN-related problems.

Usage Guidelines

Because debugging output is assigned high priority in the CPU process, it can render the system unusable. For this reason, use **debug** commands only to troubleshoot specific problems or during troubleshooting sessions with the Cisco Technical Assistance Center (TAC). Moreover, it is best to use **debug** commands during periods of lower network traffic and fewer users. Debugging during these periods decreases the likelihood that increased **debug** command processing overhead will affect system use.

You can view debug output in a CLI session only. Output is directly available when connected to the Console port, or when in the diagnostic CLI (enter **system support diagnostic-cli**). You can also view output from the regular Firepower Threat Defense CLI using the **show console-output** command.

To show debugging messages for a given feature, use the **debug** command. To disable the display of debug messages, use the **no** form of this command. Use **no debug all** to turn off all debugging commands.

debug *feature* [*subfeature*] [*level*]
no debug *feature* [*subfeature*]

Syntax Description

feature	Specifies the feature for which you want to enable debugging. To see the available features, use the **debug ?** command for CLI help.
subfeature	(Optional) Depending on the feature, you can enable debug messages for one or more subfeatures. Use ? to see the available subfeatures.
level	(Optional) Specifies the debugging level. Use ? to see the available levels.

Command Default

The default debugging level is 1.

Example

With multiple sessions running on remote access VPN, troubleshooting can be difficult, given the size of the logs. You can use the **debug webvpn condition** command to set up filters to target your debug process more precisely.

debug webvpn condition {**group** *name* | **p-ipaddress** *ip_address* [{**subnet** *subnet_mask* | **prefix** *length*}] | **reset** | **user** *name*}

Where:

- **group** *name* filters on a group policy (not a tunnel group or connection profile).
- **p-ipaddress** *ip_address* [{**subnet** *subnet_mask* | **prefix** *length*}] filters on the public IP address of the client. The subnet mask (for IPv4) or prefix (for IPv6) is optional.
- **reset** resets all filters. You can use the **no debug webvpn condition** command to turn off a specific filter.
- **user** *name* filters by username.

If you configure more than one condition, the conditions are conjoined (ANDed), so that debugs appear only if all conditions are met.

After setting up the condition filter, use the base **debug webvpn** command to turn on the debug. Setting the conditions alone does not enable the debug. Use the **show debug** and **show webvpn debug-condition** commands to view the current state of debugging.

The following shows an example of enabling a conditional debug on the user jdoe.

```
firepower# debug webvpn condition user jdoe

firepower# show webvpn debug-condition
INFO: Webvpn conditional debug is turned ON
INFO: User name filters:
INFO: jdoe

firepower# debug webvpn
INFO: debug webvpn  enabled at level 1.
```

```
firepower# show debug
debug webvpn  enabled at level 1
INFO: Webvpn conditional debug is turned ON
INFO: User name filters:
INFO: jdoe
```

Related Commands

Command	Description
show debug	Shows the currently active debug settings.
undebug	Disables debugging for a feature. This command is a synonym for **no debug**.

debug aaa

See the following commands for debugging configurations or authentication, authorization, and accounting (AAA) settings.

debug *aaa* [*accounting* | *authentication* | *authorization* | *common* | *internal* | *shim* | *url-redirect*]

Syntax Description

aaa	Enables debugging for AAA. Use ? to see the available subfeatures.
accounting	(Optional) Enables AAA accounting debugging.
authentication	(Optional) Enables AAA authentication debugging.
authorization	(Optional) Enables AAA authorization debugging.
common	(Optional) Specifies the AAA common debug level. Use ? to see the available levels.
internal	(Optional) Enables AAA internal debugging.
shim	(Optional) Specifies the AAA shim debug level. Use ? to see the available levels.
url-redirect	(Optional) Enables AAA url-redirect debugging.

Command Default

The default debugging level is 1.

Related Commands

Command	Description
show debug aaa	Shows the currently active debug settings for AAA.
undebug aaa	Disables debugging for AAA. This command is a synonym for **no debug aaa**.

debug crypto

See the following commands for debugging configurations or settings associated with crypto.

debug *crypto* [*ca* | *condition* | *engine* | *ike-common* | *ikev1* | *ikev2* | *ipsec* | *ss-apic*]

Syntax Description

crypto	Enables debugging for *crypto*. Use ? to see the available subfeatures.
ca	(Optional) Specifies the PKI debug levels. Use ? to see the available subfeatures.
condition	(Optional) Specifies the IPsec/ISAKMP debug filters. Use ? to see the available filters.
engine	(Optional) Specifies the crypto engine debug levels. Use ? to see the available levels.
ike-common	(Optional) Specifies the IKE common debug levels. Use ? to see the available levels.
ikev1	(Optional) Specifies the IKE version 1 debug levels. Use ? to see the available levels.
ikev2	(Optional) Specifies the IKE version 2 debug levels. Use ? to see the available levels.
ipsec	(Optional) Specifies the IPsec debug levels. Use ? to see the available levels.
condition	(Optional) Specifies the Crypto Secure Socket API debug levels. Use ? to see the available levels.
vpnclient	(Optional) Specifies the EasyVPN client debug levels. Use ? to see the available levels.

Command Default

The default debugging level is 1.

Related Commands

Command	Description
show debug crypto	Shows the currently active debug settings for crypto.
undebug crypto	Disables debugging for crypto. This command is a synonym for **no debug crypto**.

debug crypto ca

See the following commands for debugging configurations or settings associated with crypto ca.

debug *crypto ca* [*cluster* | *messages* | *periodic-authentication* | *scep-proxy* | *transactions* | *trustpool*] [*1-255*]

Syntax Description

crypto ca	Enables debugging for *crypto ca*. Use ? to see the available subfeatures.
cluster	(Optional) Specifies the PKI cluster debug level. Use ? to see the available levels.
cmp	(Optional) Specifies the CMP transactions debug level. Use ? to see the available levels.
messages	(Optional) Specifies the PKI Input/Output message debug level. Use ? to see the available levels.

periodic-authentication	(Optional) Specifies the PKI periodic-authentication debug level. Use ? to see the available levels.
scep-proxy	(Optional) Specifies the SCEP proxy debug level. Use ? to see the available levels.
server	(Optional) Specifies the local CA server debug level. Use ? to see the available levels.
transactions	(Optional) Specifies the PKI transaction debug level. Use ? to see the available levels.
trustpool	(Optional) Specifies the trustpool debug level. Use ? to see the available levels.
1-255	(Optional) Specifies the debugging level.

Command Default

The default debugging level is 1.

Related Commands

Command	Description
show debug crypto ca	Shows the currently active debug settings for crypto ca.
undebug	Disables debugging for crypto ca. This command is a synonym for **no debug crypto ca**.

debug crypto ikev1

See the following commands for debugging configurations or settings associated with Internet Key Exchange version 1 (IKEv1).

debug *crypto ikev1* [*timers*] [*1-255*]

Syntax Description

ikev1	Enables debugging for *ikev1*. Use ? to see the available subfeatures.
timers	(Optional) Enables debugging for IKEv1 timers.
1-255	(Optional) Specifies the debugging level.

Command Default

The default debugging level is 1.

Related Commands

Command	Description
show debug crypto ikev1	Shows the currently active debug settings for IKEv1.
undebug crypto ikev1	Disables debugging for IKEv1. This command is a synonym for **no debug crypto ikev1**.

debug crypto ikev2

See the following commands for debugging configurations or settings associated with Internet Key Exchange version 2 (IKEv2).

debug *crypto ikev2* [*ha* | *platform* | *protocol* | *timers*]

Syntax Description

ikev2	Enables debugging *ikev2*. Use ? to see the available subfeatures.
ha	(Optional) Specifies the IKEv2 HA debug level. Use ? to see the available levels.
platform	(Optional) Specifies the IKEv2 platform debug level. Use ? to see the available levels.
protocol	(Optional) Specifies the IKEv2 protocol debug level. Use ? to see the available levels.
timers	(Optional) Enables debugging for IKEv2 timers.

Command Default

The default debugging level is 1.

Related Commands

Command	Description
show debug crypto ikev2	Shows the currently active debug settings for IKEv2.
undebugcrypto ikev2	Disables debugging for IKEv2. This command is a synonym for **no debug crypto ikev2**.

debug crypto ipsec

See the following commands for debugging configurations or settings associated with IPsec.

debug *crypto ipsec* [*1-255*]

Syntax Description

ipsec	Enables debugging for *ipsec*. Use ? to see the available subfeatures.
1-255	(Optional) Specifies the debugging level.

Command Default

The default debugging level is 1.

Related Commands

Command	Description
show debug crypto ipsec	Shows the currently active debug settings for IPsec.
undebugcrypto ipsec	Disables debugging for IPsec. This command is a synonym for **no debug crypto ipsec**.

debug ldap

See the following commands for debugging configurations or settings associated with LDAP (Lightweight Directory Access Protocol).

debug *ldap* [*1-255*]

Syntax Description

ldap	Enables debugging for LDAP. Use ? to see the available subfeatures.

1-255	(Optional) Specifies the debugging level.

Command Default

The default debugging level is 1.

Related Commands

Command	Description
show debug ldap	Shows the currently active debug settings for LDAP.
undebugldap	Disables debugging for LDAP. This command is a synonym for **no debug ldap**.

debug ssl

See the following commands for debugging configurations or settings associated with SSL sessions.

debug *ssl* [*cipher* | *device*] [*1-255*]

Syntax Description

ssl	Enables debugging for SSL. Use ? to see the available subfeatures.
cipher	(Optional) Specifies the SSL cipher debug level. Use ? to see the available levels.
device	(Optional) Specifies the SSL device debug level. Use ? to see the available levels.
1-255	(Optional) Specifies the debugging level.

Command Default

The default debugging level is 1.

Related Commands

Command	Description
show debug ssl	Shows the currently active debug settings for SSL.
undebug ssl	Disables debugging for SSL. This command is a synonym for **no debug ssl**.

debug webvpn

See the following commands for debugging configurations or settings associated with WebVPN.

debug *webvpn* [*anyconnect* | *chunk* | *cifs* | *citrix* | *compression* | *condition* | *cstp-auth* | *customization* | *failover* | *html* | *javascript* | *kcd* | *listener* | *mus* | *nfs* | *request* | *response* | *saml* | *session* | *task* | *transformation* | *url* | *util* | *xml*]

Syntax Description

webvpn	Enables debugging for WebVPN. Use ? to see the available subfeatures.
anyconnect	(Optional) Specifies the WebVPN Secure Client debug level. Use ? to see the available levels.
chunk	(Optional) Specifies the WebVPN chunk debug level. Use ? to see the available levels.
cifs	(Optional) Specifies the WebVPN CIFS debug level. Use ? to see the available levels.

citrix	(Optional) Specifies the WebVPN Citrix debug level. Use ? to see the available levels.
compression	(Optional) Specifies the WebVPN compression debug level. Use ? to see the available levels.
condition	(Optional) Specifies the WebVPN filter conditions debug level. Use ? to see the available levels.
cstp-auth	(Optional) Specifies the WebVPN CSTP authentication debug level. Use ? to see the available levels.
customization	(Optional) Specifies the WebVPN customization debug level. Use ? to see the available levels.
failover	(Optional) Specifies the WebVPN failover debug level. Use ? to see the available levels.
html	(Optional) Specifies the WebVPN HTML debug level. Use ? to see the available levels.
javascript	(Optional) Specifies the WebVPN Javascript debug level. Use ? to see the available levels.
kcd	(Optional) Specifies the WebVPN KCD debug level. Use ? to see the available levels.
listener	(Optional) Specifies the WebVPN listener debug level. Use ? to see the available levels.
mus	(Optional) Specifies the WebVPN MUS debug level. Use ? to see the available levels.
nfs	(Optional) Specifies the WebVPN NFS debug level. Use ? to see the available levels.
request	(Optional) Specifies the WebVPN request debug level. Use ? to see the available levels.
response	(Optional) Specifies the WebVPN response debug level. Use ? to see the available levels.
saml	(Optional) Specifies the WebVPN SAML debug level. Use ? to see the available levels.
session	(Optional) Specifies the WebVPN session debug level. Use ? to see the available levels.
task	(Optional) Specifies the WebVPN task debug level. Use ? to see the available levels.
transformation	(Optional) Specifies the WebVPN transformation debug level. Use ? to see the available levels.

url	(Optional) Specifies the WebVPN URL debug level. Use ? to see the available levels.
util	(Optional) Specifies the WebVPN utility debug level. Use ? to see the available levels.
xml	(Optional) Specifies the WebVPN XML debug level. Use ? to see the available levels.

Command Default

The default debugging level is 1.

Related Commands

Command	Description
show debug webvpn	Shows the currently active debug settings for WebVPN.
undebug webvpn	Disables debugging for WebVPN. This command is a synonym for **no debug webvpn**.

PART VII

Access Control

- Access Control Overview, on page 1723
- Access Control Policies, on page 1745
- Access Control Rules, on page 1773
- Cisco Secure Dynamic Attributes Connector, on page 1803
- URL Filtering, on page 1857
- Security Intelligence, on page 1885
- DNS Policies, on page 1897
- Prefiltering and Prefilter Policies , on page 1915
- Service Policies, on page 1937
- Threat Detection, on page 1955
- Intelligent Application Bypass, on page 1965
- Content Restriction, on page 1973
- Zero Trust Access, on page 1979

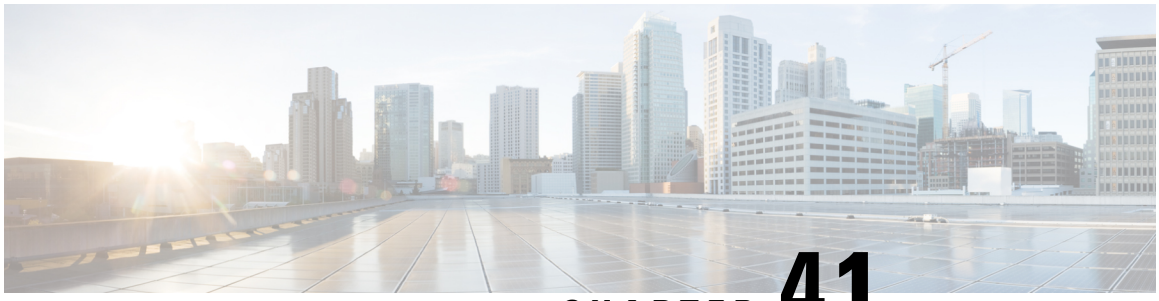

CHAPTER 41

Access Control Overview

- Introduction to Access Control, on page 1723
- Introduction to Rules, on page 1724
- Access Control Policy Default Action, on page 1726
- Deep Inspection Using File and Intrusion Policies, on page 1728
- Access Control Policy Inheritance, on page 1731
- Best Practices for Application Control, on page 1732
- Best Practices for Access Control Rules, on page 1738

Introduction to Access Control

Access control is a hierarchical policy-based feature that allows you to specify, inspect, and log (non-fast-pathed) network traffic.

Each managed device can be assigned to one access control policy. The data that the policy's assigned (also known as targeted) devices collect about your network traffic can be used to filter and control that traffic based on:

- simple, easily determined transport and network layer characteristics: source and destination, port, protocol, and so on
- the latest contextual information on the traffic, including characteristics such as reputation, risk, business relevance, application used, or URL visited
- realm, user, user group, or ISE attribute
- custom Security Group Tag (SGT)
- characteristics of encrypted traffic; you can also decrypt this traffic for further analysis
- whether unencrypted or decrypted traffic contains a prohibited file, detected malware, or intrusion attempt
- time and day (on supported devices)

Each type of traffic inspection and control occurs where it makes the most sense for maximum flexibility and performance. For example, reputation-based blocking uses simple source and destination data, so it can block prohibited traffic early in the process. In contrast, detecting and blocking intrusions and exploits is a last-line defense.

Introduction to Rules

Rules in various policy types (access control, SSL, identity, and so on) exert granular control over network traffic. The system evaluates traffic against rules in the order that you specify, using a first-match algorithm.

Although these rules might include other configurations that are not consistent across policies, they share many basic characteristics and configuration mechanics, including:

- Conditions: Rule conditions specify the traffic that each rule handles. You can configure each rule with multiple conditions. Traffic must match all conditions to match the rule.
- Action: A rule's action determines how the system handles matching traffic. Note that even if a rule does not have an **Action** list you can choose from, the rule still has an associated action. For example, a custom network analysis rule uses a network analysis policy as its "action." As another example, QoS rules do not have an explicit action because all QoS rules do the same thing: rate limit traffic.
- Position: A rule's position determines its evaluation order. When using a policy to evaluate traffic, the system matches traffic to rules in the order you specify. Usually, the system handles traffic according to the first rule where all the rule's conditions match the traffic. (Monitor rules, which are designed to track and log, are an exception.) Proper rule order reduces the resources required to process network traffic, and prevents rule preemption.
- Category: To organize some rule types, you can create custom rule categories in each parent policy.
- Logging: For many rules, logging settings govern whether and how the system logs connections handled by the rule. Some rules (such as identity and network analysis rules) do not include logging settings because the rules neither determine the final disposition of connections, nor are they specifically designed to log connections. As another example, QoS rules do not include logging settings; you cannot log a connection simply because it was rate limited.
- Comments: For some rule types, each time you save changes, you can add comments. For example, you might summarize the overall configuration for the benefit of other users, or note when you change a rule and the reason for the change.

Tip A right-click menu in many policy editors provides shortcuts to many rule management options, including editing, deleting, moving, enabling, and disabling.

For more information, see the chapter that discusses the rules you're interested in (for example, access control rules).

Related Topics

Configuring Application Conditions and Filters, on page 1789
Best Practices for Application Control, on page 1732

Filtering Rules by Device

Some policy editors allow you to filter your rule view by affected devices.

The system uses a rule's interface constraints to determine if the rule affects a device. If you constrain a rule by interface (security zone or interface group condition), the device where that interface is located is affected by that rule. Rules with no interface constraint apply to any interface, and therefore every device.

QoS rules are always constrained by interface.

Note The following procedure does not apply to the access control policy. To see the rules that apply to a specific device or set of devices in the access control policy, click the filter icon and select the devices.

Procedure

Step 1 In the policy editor, click **Rules**, then click **Filter by Device**.
A list of targeted devices and device groups appears.

Step 2 Check one or more check boxes to display only the rules that apply to those devices or groups. Or, check **All** to reset and display all of the rules.

Tip
Hover your pointer over a rule criterion to see its value. If the criterion represents an object with device-specific overrides, the system displays the override value when you filter the rules list by only that device. If the criterion represents an object with domain-specific overrides, the system displays the override value when you filter the rules list by devices in that domain.

Step 3 Click **OK**.

Rule and Other Policy Warnings

Policy and rule editors use icons to mark configurations that could adversely affect traffic analysis and flow. Depending on the issue, the system may warn you when you deploy or prevent you from deploying entirely.

Tip Hover your pointer over an icon to read the warning, error, or informational text.

Table 92: Policy Error Icons

Icon	Description	Example
Errors (✖)	If a rule or configuration has an error, you cannot deploy until you correct the issue, even if you disable any affected rules.	A rule that performs category and reputation-based URL filtering is valid until you assign a device that does not have a URL Filtering license. At that point, an error icon appears next to the rule, and you cannot deploy until you edit or delete the rule, unassign the device, or enable the license.

Icon	Description	Example
Warning (⚠)	You can deploy a policy that displays rule or other warnings. However, misconfigurations marked with warnings have no effect. If you disable a rule with a warning, the warning icon disappears. It reappears if you enable the rule without correcting the underlying issue.	Preempted rules or rules that cannot match traffic due to misconfiguration have no effect. This includes conditions using empty object groups, application filters that match no applications, excluded LDAP users, invalid ports, and so on. However, if a warning icon marks a licensing error or model mismatch, you cannot deploy until you correct the issue.
Information (ⓘ)	Information icons convey helpful information about configurations that may affect the flow of traffic. These issues do not prevent you from deploying.	The system might skip matching the first few packets of a connection against some rules, until the system identifies the application or web traffic in that connection. This allows connections to be established so that applications and HTTP requests can be identified.
Rule Conflict (∿)	When you enable rule conflict analysis, this icon appears in the rule table for rules that have conflicts.	Conflicts include redundant rules, redundant objects, and shadowed rules. Redundant and shadowed rules do not match traffic because previous rules would already match the criteria. Redundant objects make your rules unnecessarily complex.

Access Control Policy Default Action

A newly created access control policy directs its assigned devices to handle all traffic using its *default action*.

In a simple access control policy, the default action specifies how a device handles all traffic. In a more complex policy, the default action handles traffic that:

- is not trusted by Intelligent Application Bypass
- is not on a Security Intelligence Block list
- is not blocked by SSL inspection (encrypted traffic only)
- matches none of the rules in the policy (except Monitor rules, which match and log—but do not handle or inspect—traffic)

The access control policy default action can block or trust traffic without further inspection, or inspect traffic for intrusions and discovery data.

Note You **cannot** perform file or malware inspection on traffic handled by the default action. Logging for connections handled by the default action is initially disabled, though you can enable it.

If you are using policy inheritance, the default action for the lowest-level descendant determines final traffic handling. Although an access control policy can inherit its default action from its base policy, you cannot enforce this inheritance.

The following table describes the types of inspection you can perform on traffic handled by each default action.

Table 93: Access Control Policy Default Actions

Default Action	Effect on Traffic	Inspection Type and Policy
Access Control: Block All Traffic	block without further inspection	none
Access Control: Trust All Traffic	trust (allow to its final destination without further inspection)	none
Intrusion Prevention	allow, as long as it is passed by the intrusion policy you specify	intrusion, using the specified intrusion policy and associated variable set, and discovery, using the network discovery policy
Network Discovery Only	allow	discovery only, using the network discovery policy
Inherit from base policy	defined in base policy	defined in base policy

The following diagram illustrates the table.

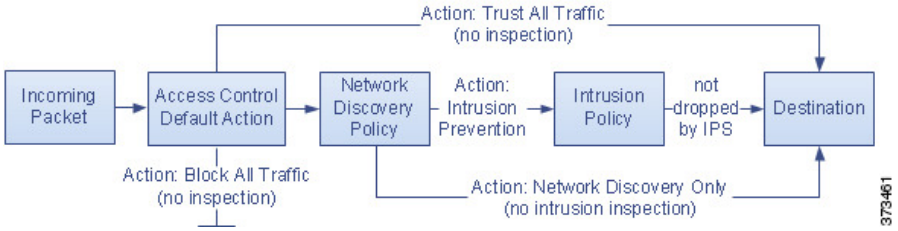

The following diagrams illustrate the **Block All Traffic** and **Trust All Traffic** default actions.

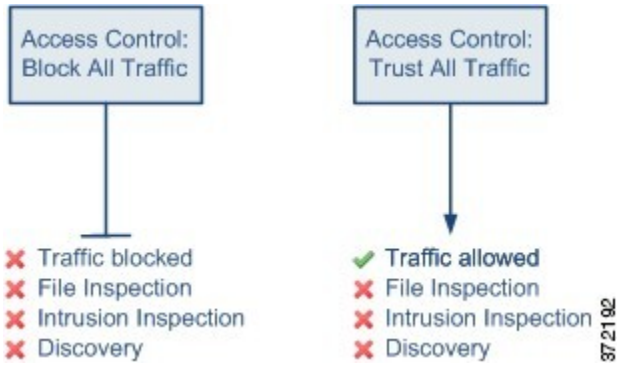

The following diagrams illustrate the **Intrusion Prevention** and **Network Discovery Only** default actions.

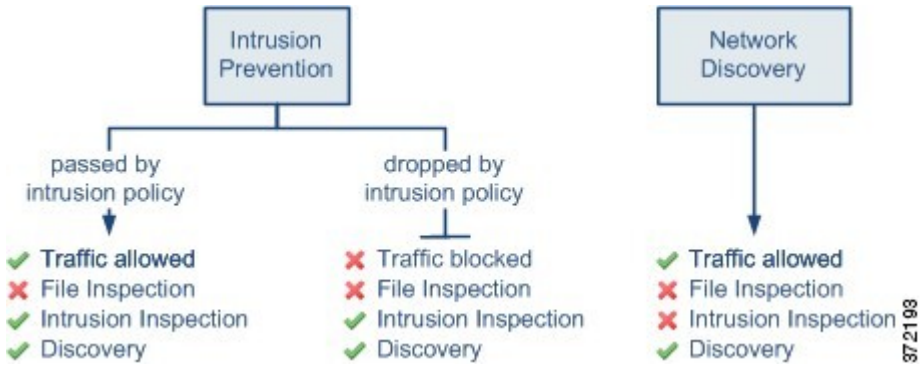

Tip The purpose of **Network Discovery Only** is to improve performance in a discovery-only deployment. Different configurations can disable discovery if you are only interested in intrusion detection and prevention.

Deep Inspection Using File and Intrusion Policies

Deep inspection uses intrusion and file policies as the last line of defense before traffic is allowed to its destination.

- *Intrusion policies* govern the system's intrusion prevention capabilities.
- *File policies* govern the system's file control and malware defense capabilities.

  For complete information, see Network Malware Protection and File Policies, on page 2209.

Access control occurs before deep inspection; access control rules and the access control default action determine which traffic is inspected by intrusion and file policies.

By associating an intrusion or file policy with an access control rule, you are telling the system that before it passes traffic that matches the access control rule's conditions, you first want to inspect the traffic with an intrusion policy, a file policy, or both.

In an access control policy, you can associate one intrusion policy with each Allow and Interactive Block rule, as well as with the default action. Every unique **pair** of intrusion policy and variable set counts as one policy.

To associate intrusion and file policies with an access control rule, see:

- Access Control Rule Configuration to Perform Intrusion Prevention
- Configuring an Access Control Rule to Perform Malware Protection, on page 2217

Note By default, the system disables intrusion and file inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion and file inspection configured.

Related Topics

How Policies Examine Traffic For Intrusions
File Policies, on page 2210

Access Control Traffic Handling with Intrusion and File Policies

The following diagram shows the flow of traffic in an inline intrusion prevention and malware defense deployment, as governed by an access control policy that contains four different types of access control rules and a default action.

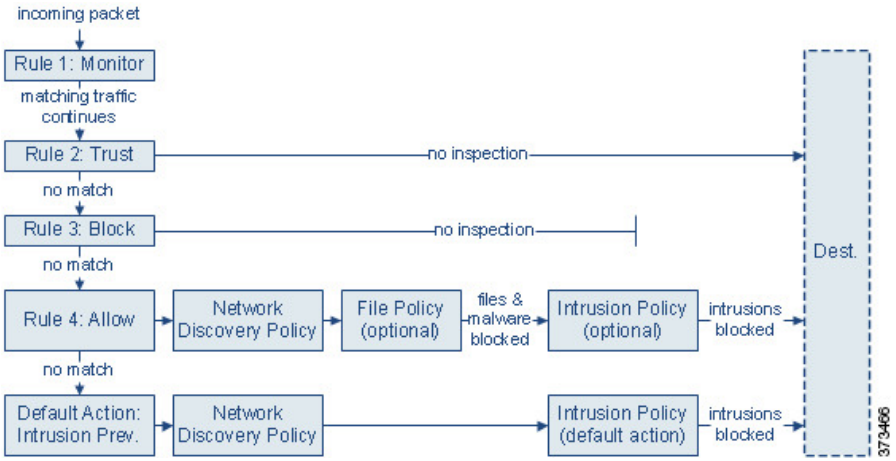

In the scenario above, the first three access control rules in the policy—Monitor, Trust, and Block—cannot inspect matching traffic. Monitor rules track and log but do not inspect network traffic, so the system continues to match traffic against additional rules to determine whether to permit or deny it. (However, see an important exception and caveat at Access Control Rule Monitor Action, on page 1778.) Trust and Block rules handle matching traffic without further inspection of any kind, while traffic that does not match continues to the next access control rule.

The fourth and final rule in the policy, an Allow rule, invokes various other policies to inspect and handle matching traffic, in the following order:

- **Discovery: Network Discovery Policy**—First, the network discovery policy inspects traffic for discovery data. Discovery is passive analysis and does not affect the flow of traffic. Although you do not explicitly enable discovery, you can enhance or disable it. However, allowing traffic does not automatically guarantee discovery data collection. The system performs discovery only for connections involving IP addresses that are explicitly monitored by your network discovery policy.
- **malware defense and File Control: File Policy**—After traffic is inspected by discovery, the system can inspect it for prohibited files and malware. malware defense detects and optionally blocks malware in many types of files, including PDFs, Microsoft Office documents, and others. If your organization wants to block not only the transmission of malware files, but all files of a specific type (regardless of whether the files contain malware), *file control* allows you to monitor network traffic for transmissions of specific file types, then either block or allow the file.
- **Intrusion Prevention: Intrusion Policy**—After file inspection, the system can inspect traffic for intrusions and exploits. An intrusion policy examines decoded packets for attacks based on patterns, and can block or alter malicious traffic. Intrusion policies are paired with *variable sets*, which allow you to use named values to accurately reflect your network environment.

- **Destination**—Traffic that passes all the checks described above passes to its destination.

An Interactive Block rule (not shown in the diagram) has the same inspection options as an Allow rule. This is so you can inspect traffic for malicious content when a user bypasses a blocked website by clicking through a warning page.

Traffic that does not match any access control rules in the policy with an action other than Monitor is handled by the default action. In this scenario, the default action is an Intrusion Prevention action, which allows traffic to its final destination as long as it is passed by the intrusion policy you specify. In a different deployment, you might have a default action that trusts or blocks all traffic without further inspection. Note that the system can inspect traffic allowed by the default action for discovery data and intrusions, but not prohibited files or malware. You **cannot** associate a file policy with the access control default action.

Note Sometimes, when a connection is analyzed by an access control policy, the system must process the first few packets in that connection, **allowing them to pass**, before it can decide which access control rule (if any) will handle the traffic. However, so these packets do not reach their destination uninspected, you can specify an intrusion policy (in the Advanced settings for the access control policy) to inspect these packets and generate intrusion events.

File and Intrusion Inspection Order

In your access control policy, you can associate multiple Allow and Interactive Block rules with different intrusion and file policies to match inspection profiles to various types of traffic.

Note Traffic allowed by an Intrusion Prevention or Network Discovery Only default action can be inspected for discovery data and intrusions, but cannot be inspected for prohibited files or malware. You **cannot** associate a file policy with the access control default action.

You do not have to perform both file and intrusion inspection in the same rule. For a connection matching an Allow or Interactive Block rule:

- without a file policy, traffic flow is determined by the intrusion policy
- without an intrusion policy, traffic flow is determined by the file policy
- without either, allowed traffic is inspected by network discovery only

Tip The system does not perform any kind of inspection on trusted traffic. Although configuring an Allow rule with neither an intrusion nor file policy passes traffic like a Trust rule, Allow rules let you perform discovery on matching traffic.

The diagram below illustrates the types of inspection you can perform on traffic that meets the conditions of either an Allow or user-bypassed Interactive Block access control rule. For simplicity, the diagram displays traffic flow for situations where both (or neither) an intrusion and a file policy are associated with a single access control rule.

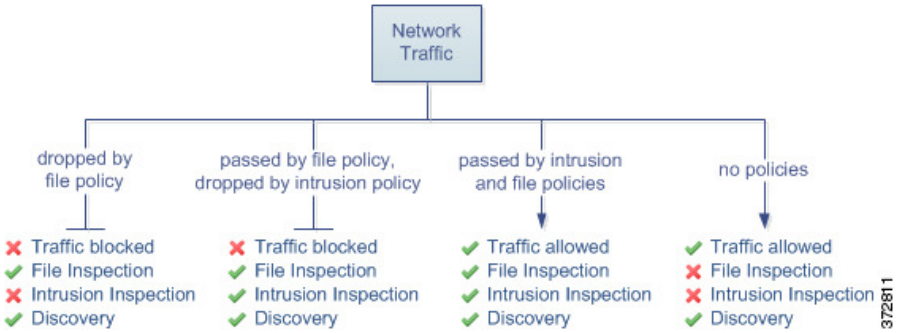

For any single connection handled by an access control rule, file inspection occurs before intrusion inspection. That is, the system does not inspect files blocked by a file policy for intrusions. Within file inspection, simple blocking by type takes precedence over malware inspection and blocking.

For example, consider a scenario where you normally want to allow certain network traffic as defined in an access control rule. However, as a precaution, you want to block the download of executable files, examine downloaded PDFs for malware and block any instances you find, and perform intrusion inspection on the traffic.

You create an access control policy with a rule that matches the characteristics of the traffic you want to provisionally allow, and associate it with both an intrusion policy and a file policy. The file policy blocks the download of all executables, and also inspects and blocks PDFs containing malware:

- First, the system blocks the download of all executables, based on simple type matching specified in the file policy. Because they are immediately blocked, these files are subject to neither malware nor intrusion inspection.
- Next, the system performs malware cloud lookups for PDFs downloaded to a host on your network. Any PDFs with a malware disposition are blocked, and are not subject to intrusion inspection.
- Finally, the system uses the intrusion policy associated with the access control rule to inspect any remaining traffic, including files not blocked by the file policy.

Note Until a file is detected and blocked in a session, packets from the session may be subject to intrusion inspection.

Access Control Policy Inheritance

You can nest access control policies, where each policy inherits the rules and settings from an ancestor (or *base*) policy. You can enforce this inheritance, or allow lower-level policies to override their ancestors.

Access control uses a hierarchical policy-based implementation. Just as you create a domain hierarchy, you can create a corresponding hierarchy of access control policies. A *descendant*, or *child*, access control policy inherits rules and settings from its direct *parent*, or base, policy. That base policy may have its own parent policy from which it inherits rules and settings, and so on.

An access control policy's rules are nested between its parent policy's Mandatory and Default rule sections. This implementation enforces Mandatory rules from ancestor policies, but allows the current policy to write rules that preempt Default rules from ancestor policies.

You can lock the following settings to enforce them in all descendant policies. Descendant policies can override unlocked settings.

- Security Intelligence — connections that are allowed or blocked based on the latest reputation intelligence for IP addresses, URLs, and domain names.
- HTTP Response pages — Displaying a custom or system-provided response page when you block a user's website request.
- Advanced settings — Specifying associated subpolicies, network analysis settings, performance settings, and other general options.

When using policy inheritance, the default action for the lowest-level descendant determines final traffic handling. Although an access control policy can inherit its default action from an ancestor policy, you cannot enforce this inheritance.

Policy Inheritance and Multitenancy

Access control's hierarchical policy-based implementation complements multitenancy.

In a typical multidomain deployment, access control policy hierarchy corresponds to domain structure, and you apply the lowest-level access control policy to managed devices. This implementation allows selective access control enforcement at a higher domain level, while lower-level domain administrators can tailor deployment-specific settings. (You must use roles, not policy inheritance and enforcement alone, to restrict administrators in descendant domains.)

For example, as a Global domain administrator for your organization, you can create an access control policy at the Global level. You can then require that all your devices, which are divided into subdomain by function, use that Global-level policy as a base policy.

When subdomain administrators log into the Secure Firewall Management Center to configure access control, they can deploy the Global-level policy as-is. Or, they can create and deploy a descendant access control policy within the boundaries of the Global-level policy.

Note Although the most useful implementation of access control inheritance and enforcement complements multitenancy, you can create a hierarchy of access control policies within a single domain. You can also assign and deploy access control policies at any level.

Best Practices for Application Control

The following topics discuss our recommended best practices for controlling applications with access control rules.

Recommendations for Application Control

Keep in mind the following guidelines and limitations for application control:

Ensuring that Adaptive Profiling is Enabled

If adaptive profiling is not enabled (its default state), access control rules cannot perform application control.

Automatically Enabling Application Detectors

If no detector is enabled for an application you want to detect, the system automatically enables all system-provided detectors for the application. If none exist, the system enables the most recently modified user-defined detector for the application.

Configure Your Policy to Examine the Packets That Must Pass Before an Application Is Identified

The system cannot perform application control, including Intelligent Application Bypass (IAB) and rate limiting, before *both* of the following occur:

- A monitored connection is established between a client and server
- The system identifies the application in the session

This identification should occur in 3 to 5 packets, or after the server certificate exchange in the SSL handshake if the traffic is encrypted.

If early traffic matches all other criteria but application identification is incomplete, the system allows the packet to pass and the connection to be established (or the SSL handshake to complete). After the system completes its identification, the system applies the appropriate action to the remaining session traffic.

Note A server must adhere to the protocol requirements of an application for the system to be able to recognize it. For example, if you have a server that sends a keep-alive packet rather than an ACK when an ACK is expected, that application might not be identified, and the connection will not match the application-based rule. Instead, it will be handled by another matching rule or the default action. This might mean that connections you want to allow can be denied instead. If you run into this problem, and you cannot fix the server to follow the protocol standards, you need to write a non-application-based rule to cover traffic for that server, for example, by matching the IP address and port number.

Create Separate Rules for URL and Application Filtering

Create separate rules for URL and application filtering whenever possible, because combining application and URL criteria can lead to unexpected results, especially for encrypted traffic.

Rules that include both application and URL criteria should come after application-only or URL-only rules, unless the application+URL rule is acting as an exception to a more general application-only or URL-only rule.

URL Rules Before Application and Other Rules

For the most effective URL matching, place rules that include URL conditions before other rules, particularly if the URL rules are block rules and the other rules meet both of the following criteria:

- They include application conditions.
- The traffic to be inspected is encrypted.

Application Control for Encrypted and Decrypted Traffic

The system can identify and filter encrypted and decrypted traffic:

- Encrypted traffic—The system can detect application traffic encrypted with StartTLS, including SMTPS, POPS, FTPS, TelnetS, and IMAPS. In addition, it can identify certain encrypted applications based on the Server Name Indication in the TLS ClientHello message, or the subject distinguished name value from the server certificate. These applications are tagged `SSL Protocol`; in an SSL rule, you can choose only these applications. Applications without this tag can only be detected in unencrypted or decrypted traffic.
- Decrypted traffic—The system assigns the `decrypted traffic` tag to applications that the system can detect in decrypted traffic only, not encrypted or unencrypted.

TLS Server Identity Discovery and Application Control

The latest version of the Transport Layer Security (TLS) protocol 1.3, defined by RFC 8446, is the preferred protocol for many web servers to provide secure communications. Because the TLS 1.3 protocol encrypts the server's certificate for additional security, and the certificate is needed to match application and URL filtering criteria in access control rules, the Firepower System provides a way to extract the server certificate *without* decrypting the entire packet.

We strongly recommend enabling it for any traffic you want to match on application or URL criteria, especially if you want to perform deep inspection of that traffic. A decryption policy is not required because *traffic is not decrypted* in the process of extracting the server certificate.

For more information, see Access Control Policy Advanced Settings, on page 1756.

Exempting Applications from Active Authorization

In an identity policy, you can exempt certain applications from active authentication, allowing traffic to continue to access control. These applications are tagged `User-Agent Exclusion`. In an identity rule, you can choose only these applications.

Handling Application Traffic Packets Without Payloads

When performing access control, the system applies the default policy action to packets that do not have a payload in a connection where an application is identified.

Handling Referred Application Traffic

To handle traffic referred by a web server, such as advertisement traffic, match the referred application rather than the referring application.

Controlling Application Traffic That Uses Multiple Protocols (Skype, Zoho)

Some applications use multiple protocols. To control their traffic, make sure your access control policy covers all relevant options. For example:

- Skype—To control Skype traffic, choose the **Skype** tag from the **Application Filters** list rather than selecting individual applications. This ensures that the system can detect and control all Skype traffic the same way.
- Zoho—To control Zoho mail, choose *both* **Zoho** and **Zoho mail** from the Available Application list.

Search Engines Supported for Content Restriction Features

The system supports Safe Search filtering for specific search engines only. The system assigns the `safesearch supported` tag to application traffic from these search engines.

Controlling Evasive Application Traffic

See Application-Specific Notes and Limitations, on page 1737.

Best Practices for Configuring Application Control

We recommend controlling applications' access to the network as follows:

- To allow or block application access from a less secure network to a more secure network: Use **Port** (Selected Destination Port) conditions on the access control rule

  For example, allow ICMP traffic from the internet (less secure) to an internal network (more secure.)

- To allow or block applications being accessed by user groups: Use **Application** conditions on the access control rule

  For example, block Facebook from being accessed by members of the Contractors group

Caution Failure to set up your access control rules properly can have unexpected results, including traffic being allowed that should be blocked. In general, application control rules should be lower in your access control list because it takes longer for those rules to match than rules based on IP address, for example.

Access control rules that use *specific* conditions (such as networks and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your access control rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your access control rules. For more information about the OSI model, see this Wikipedia article.

The following table provides an example of how to set up your access control rules:

Type of control	Action	Zones, Networks, VLAN Tags	Users	Applications	Ports	URLs	SGT/ISE Attributes	Inspection, Logging, Comments
Application from more secure to less secure network when application uses a port (for example, SSH)	Your choice (**Allow** in this example)	Destination zones or networks using the outside interface	Any	Do not set	Available Ports : **SSH** Add to **Selected Destination Ports**	Any	Use only with ISE/ISE-PIC.	Any

Type of control	Action	Zones, Networks, VLAN Tags	Users	Applications	Ports	URLs	SGT/ISE Attributes	Inspection, Logging, Comments
Application from more secure to less secure network when application does not use a port (for example, ICMP)	Your choice (**Allow** in this example)	Destination zones or networks using the outside interface	Any	Do not set	Selected Destination Ports **Protocol**: `ICMP` **Type**: `Any`	Do not set	Use only with ISE/ISE-PIC.	Any
Application access by a user group	Your choice (**Block** in this example)	Your choice	Choose a user group (Contractors group in this example)	Choose the name of the application (**Facebook** in this example)	Do not set	Do not set	Use only with ISE/ISE-PIC.	Your choice

Application Characteristics

The system characterizes each application that it detects using the criteria described in the following table. Use these characteristics as application filters.

Table 94: Application Characteristics

Characteristic	Description	Example
Type	Application protocols represent communications between hosts. Clients represent software running on a host. Web applications represent the content or requested URL for HTTP traffic.	HTTP and SSH are application protocols. Web browsers and email clients are clients. MPEG video and Facebook are web applications.
Risk	The likelihood that the application is being used for purposes that might be against your organization's security policy.	Peer-to-peer applications tend to have a very high risk.
Business Relevance	The likelihood that the application is being used within the context of your organization's business operations, as opposed to recreationally.	Gaming applications tend to have a very low business relevance.
Category	A general classification for the application that describes its most essential function. Each application belongs to at least one category.	Facebook is in the social networking category.

Characteristic	Description	Example
Tag	Additional information about the application. Applications can have any number of tags, including none.	Video streaming web applications often are tagged `high bandwidth` and `displays ads`.

Application-Specific Notes and Limitations

- Office 365 Admin Portal:

  Limitation: If the access policy has logging enabled at the beginning as well as at the end, the first packet will be detected as Office 365 and the end of connection will be detected as Office 365 Admin Portal. This should not affect blocking.

- Skype:

  See Recommendations for Application Control, on page 1732

- GoToMeeting

  In order to fully detect GoToMeeting, your rule must include all of the following applications:

  - GoToMeeting
  - Citrix Online
  - Citrix GoToMeeting Platform
  - LogMeIn
  - STUN

- Zoho:

  See Recommendations for Application Control, on page 1732

- Evasive applications such as Bittorrent, Tor, Psiphon, and Ultrasurf:

  For evasive applications, only the highest-confidence scenarios are detected by default. If you need to take action on this traffic (such as block or implement QoS), it may be necessary to configure more aggressive detection with better effectiveness. To do this, contact TAC to review your configurations as these changes may result in false positives.

- WeChat:

  It is not possible to selectively block WeChat Media if you allow WeChat.

- RDP (Remote Desktop Protocol):

  If allowing the RDP application does not allow file transfers, ensure that the rule for RDP includes both the TCP and UDP port 3389. RDP file transfer uses UDP.

Best Practices for Access Control Rules

Properly configuring and ordering rules is essential to building an effective deployment. The following topics summarize rule performance guidelines.

Note When you deploy configuration changes, the system evaluates all rules together and creates an expanded set of criteria that assigned devices use to evaluate network traffic. If these criteria exceed the resources (physical memory, processors, and so on) of a device, you cannot deploy to that device.

General Best Practices for Access Control

Review the following requirements and general best practices:

- Use a prefilter policy to provide early blocking for unwanted traffic, and to fastpath traffic that does not benefit from access control inspection. For more information, see Best Practices for Fastpath Prefiltering, on page 1920.
- Although you can configure the system without licensing your deployment, many features require that you enable the appropriate licenses before you deploy.
- Access control rules are deployed as access control lists on the device. To minimize the number of access control entries created per access control rule, and improve overall performance, enable object group search for each device. Object group search is a device setting, not an access control policy setting, so you must edit each device to ensure the feature is enabled. For more information, see Configure Object Group Search, on page 199.
- When you deploy an access control policy, its rules are not applied to existing connections. Traffic on an existing connection is not bound by the new policy that is deployed. In addition, the policy hit count is incremented only for the first packet of a connection that matches a policy. Thus, the traffic on an existing connection that could match a policy is omitted from the hit count. To have the policy rules effectively applied, clear the existing connections sessions, and then deploy the policy.
- Whenever possible, combine multiple network objects into a single object group. The system automatically creates an object group (during deployment) when you select more than one object (for source or destination separately). Selecting existing groups can avoid object group duplication and reduce the potential impact on CPU usage when there are a large number of duplicate objects.
- For the system to affect traffic, you must deploy relevant configurations to managed devices using routed, switched, or transparent interfaces, or inline interface pairs.

  Sometimes, the system prevents you from deploying inline configurations to passively deployed devices, including inline devices in tap mode.

  In other cases, the policy may deploy successfully, but attempting to block or alter traffic using passively deployed devices can have unexpected results. For example, the system may report multiple beginning-of-connection events for each blocked connection, because blocked connections are not blocked in passive deployments.
- Certain features, including URL filtering, application detection, rate limiting, and Intelligent Application Bypass, must allow some packets to pass in order for the system to identify the traffic.

- You cannot perform file or malware inspection on traffic handled by the access control policy's default action.
- Some features are only available on certain device models. Warning icons and confirmation dialog boxes designate unsupported features.
- If you will use syslog or store events externally, avoid special characters in object names such as policy and rule names. Object names should not contain special characters, such as commas, that the receiving application may use as separators.
- Logging for connections handled by the default action is initially disabled, though you can enable it.
- Best practices for creating, ordering, and implementing access control rules are detailed in Best Practices for Access Control Rules, on page 1738 and subtopics.

Best Practices for Ordering Rules

General guidelines:

- In general, place top-priority rules that must apply to all traffic near the top of the policy.
- Specific rules should come before general rules, especially when the specific rules are exceptions to general rules.

  Otherwise, traffic will match the general rule first and never hit the applicable specific rule.
- Rules that drop traffic based on layer-3/4 criteria only (such as IP address, security zone, and port number) should come as early as possible. Rules based on these criteria do not require inspection to identify matching connections.
- Whenever possible, put specific drop rules near the top of the policy. This ensures the earliest possible decision on undesirable traffic.
- URL filtering, application-based, and geolocation-based rules and others that require inspection should come after rules that drop traffic based on layer-3/4 criteria only (such as IP address, security zone, and port number), but before rules that specify file and intrusion policies.
- Put URL filtering rules above application rules, and follow application rules with micro-application rules and Common Industrial Protocol (CIP) sub-classification application filtering rules.
- Rules that specify file policies and intrusion policies should come at the bottom of the rule order. These rules require resource-intensive deep inspection, and you should eliminate as many threats as possible using less-intensive methods first, for performance reasons, in order to minimize the number of potential threats that require deep inspection.
- Always order rules to suit your organization's needs.

Exceptions and additions to the above guidelines are noted in the sections below.

Rule Preemption

Rule preemption occurs when a rule will never match traffic because a rule earlier in the evaluation order matches the traffic first. A rule's conditions govern whether it preempts other rules. In the following example, the second rule cannot block Admin traffic because the first rule allows it:

Access Control Rule 1: allow Admin users

Access Control Rule 2: block Admin users

Any type of rule condition can preempt a subsequent rule. The VLAN range in the first SSL rule includes the VLAN in the second rule, so the first rule preempts the second:

SSL Rule 1: do not decrypt VLAN 22-33
SSL Rule 2: block VLAN 27

In the following example, Rule 1 matches any VLAN because no VLANs are configured, so Rule 1 preempts Rule 2, which attempts to match VLAN 2:

Access Control Rule 1: allow Source Network 10.4.0.0/16
Access Control Rule 2: allow Source Network 10.4.0.0/16, VLAN 2

A rule also preempts an identical subsequent rule where all configured conditions are the same:

QoS Rule 1: rate limit VLAN 1 URL www.netflix.com
QoS Rule 2: rate limit VLAN 1 URL www.netflix.com

A subsequent rule would not be preempted if any condition is different:

QoS Rule 1: rate limit VLAN 1 URL www.netflix.com
QoS Rule 2: rate limit VLAN 2 URL www.netflix.com

Example: Ordering SSL Rules to Avoid Preemption

Consider a scenario where a trusted CA (Good CA) mistakenly issued a CA certificate to a malicious entity (Bad CA), but has not yet revoked that certificate. You want to use an SSL policy to block traffic encrypted with certificates issued by the untrusted CA, but otherwise allow traffic within the trusted CA's chain of trust. After you upload the CA certificates and all intermediate CA certificates, configure an SSL policy with rules in the following order:

SSL Rule 1: Block issuer CN=www.badca.com
SSL Rule 2: Do not decrypt issuer CN=www.goodca.com

If you reverse the rules, you first match all traffic trusted by Good CA, including traffic trusted by Bad CA. Because no traffic ever matches the subsequent Bad CA rule, malicious traffic may be allowed instead of blocked.

Rule Actions and Rule Order

A rule's action determines how the system handles matching traffic. Improve performance by placing rules that do not perform or ensure further traffic handling before the resource-intensive rules that do. Then, the system can divert traffic that it might otherwise have inspected.

The following examples show how you might order rules in various policies, given a set of rules where none is more critical and preemption is not an issue.

If your rules include application conditions, also see Best Practices for Configuring Application Control, on page 1735.

Optimum Order: Decryption Rules

Not only does decryption require resources, but so does further analysis of the decrypted traffic. Place rules that decrypt traffic last.

Note Certain managed devices support encrypting and decrypting TLS/SSL traffic in hardware, which significantly improves performance. For more information, see TLS Crypto Acceleration, on page 2268.

1. Monitor—Rules that log matching connections, but take no other action on traffic.
2. Block, Block with reset—Rules that block traffic without further inspection.
3. Do not decrypt—Rules that do not decrypt encrypted traffic, passing the encrypted session to access control rules. The payloads of these sessions are not subject to deep inspection.
4. Decrypt - Known Key—Rules that decrypt incoming traffic with a known private key.
5. Decrypt - Resign—Rules that decrypt outgoing traffic by re-signing the server certificate.

Optimum Order: Access Control Rules

Intrusion, file, and malware inspection requires resources, especially if you use multiple custom intrusion policies and variable sets. Place access control rules that invoke deep inspection last.

1. Monitor—Rules that log matching connections, but take no other action on traffic. (However, see the important exception and caveat at Access Control Rule Monitor Action, on page 1778.)
2. Trust, Block, Block with reset—Rules that handle traffic without further inspection. Note that trusted traffic is subject to authentication requirements imposed by an identity policy, and to rate limiting.
3. Allow, Interactive Block (no deep inspection)—Rules that do not inspect traffic further, but allow discovery. Note that allowed traffic is subject to authentication requirements imposed by an identity policy, and to rate limiting.
4. Allow, Interactive Block (deep inspection)—Rules associated with file or intrusion policies that perform deep inspection for prohibited files, malware, and exploits.

Application Rule Order

Rules with application conditions are more likely to match traffic if you move them to a lower order in your list of rules.

Access control rules that use *specific* conditions (such as networks and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your access control rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your access control rules. For more information about the OSI model, see this Wikipedia article.

For more information and an example, see Best Practices for Configuring Application Control, on page 1735 and Recommendations for Application Control, on page 1732.

URL Rule Order

For the most effective URL matching, place rules that include URL conditions before other rules, particularly if the URL rules are block rules and the other rules meet both of the following criteria:

- They include application conditions.

- The traffic to be inspected is encrypted.

If you configure exceptions to a rule, put the exception above the other rule.

Best Practices for Simplifying and Focusing Rules

Simplify: Do Not Overconfigure

Minimize individual rule criteria. Use as few individual elements in rule conditions as possible. For example, in network conditions use IP address blocks rather than individual IP addresses.

If one condition is enough to match the traffic you want to handle, do not use two. Using conditions that are redundant can greatly expand the deployed configuration, which can lead to problems in device performance, and unexpected device behavior in a cluster and high-availability unit re-join. For example:

- Use security zones that represent multiple interfaces carefully. If you specify source and destination networks as conditions, and these are enough to match the traffic you are targeting, then specifying a security zone is not required.
- If you want to match a set of internal interfaces to ANY destination on the Internet (for example), then simply use a source security zone that includes your internal interfaces. No network or destination interface criteria are needed.

Combining elements into objects does **not** improve performance. For example, using a network object that contains 50 individual IP addresses gives you only an organizational—not a performance—benefit over including those IP addresses in the condition individually.

For recommendations related to application detection, see Best Practices for Configuring Application Control, on page 1735.

Focus: Narrowly Constrain Resource-Intensive Rules, Especially by Interface

As much as possible, use rule conditions to narrowly define the traffic handled by resource-intensive rules. Focused rules are also important because rules with broad conditions can match many different types of traffic, and can preempt later, more specific rules. Examples of resource-intensive rules include:

- TLS/SSL rules that decrypt traffic—Not only the decryption, but further analysis of the decrypted traffic, requires resources. Narrow focus, and where possible, block or choose not to decrypt encrypted traffic.

  Certain Firewall Threat Defense models perform TLS/SSL encryption and decryption in hardware, which improves performance significantly. For more information, see TLS Crypto Acceleration, on page 2268.
- Access control rules that invoke deep inspection—Intrusion, file, and malware inspection requires resources, especially if you use multiple custom intrusion policies and variable sets. Make sure you only invoke deep inspection where required.

For maximum performance benefit, constrain rules by interface. If a rule excludes all of a device's interfaces, that rule does not affect that device's performance.

Maximum Number of Access Control Rules and Intrusion Policies

The maximum number of access control rules or intrusion policies that are supported by a device depends on many factors, including policy complexity, physical memory, and the number of processors on the device.

If you exceed the maximum supported by a device, you cannot deploy the access control policy and you must reevaluate.

Guidelines for intrusion policies:

- In an access control policy, you can associate one intrusion policy with each Allow and Interactive Block rule, as well as with the default action. Every unique **pair** of intrusion policy and variable set counts as one policy.
- You may want to consolidate intrusion policies or variable sets so you can associate a single intrusion policy-variable set pair with multiple access control rules. On some devices you may find you can use only a single variable set for all your intrusion policies, or even a single intrusion policy-variable set pair for the whole device.

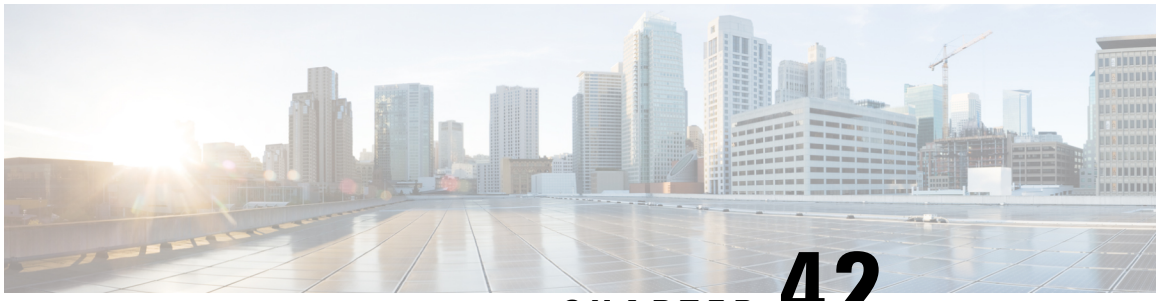

CHAPTER 42

Access Control Policies

The following topics describe how to work with access control policies:

- Access Control Policy Components, on page 1745
- System-Created Access Control Policies, on page 1746
- Requirements and Prerequisites for Access Control Policies, on page 1746
- Managing Access Control Policies, on page 1747
- History for Access Control Policies, on page 1769

Access Control Policy Components

Following are the main elements of an access control policy.

Name and Description

Each access control policy must have a unique name. A description is optional.

Inheritance Settings

Policy inheritance allows you to create a hierarchy of access control policies. A parent (or *base*) policy defines and enforces default settings for its descendants.

A policy's inheritance settings allow you to select its base policy. You can also lock settings in the current policy to force any descendants to inherit them. Descendant policies can override unlocked settings.

Policy Assignment

Each access control policy identifies the devices that use it. Each device can be targeted by only one access control policy. You can also assign the policy to device templates.

Rules

Access control rules provide a granular method of handling network traffic. Rules in an access control policy are numbered, starting at 1, including rules inherited from ancestor policies. The system matches traffic to access control rules in top-down order by ascending rule number.

Usually, the system handles network traffic according to the *first* access control rule where *all* the rule's conditions match the traffic. Conditions can be simple or complex, and their use often depends on certain licenses.

Default Action

The default action determines how the system handles and logs traffic that is not handled by any other access control configuration. The default action can block or trust all traffic without further inspection, or inspect traffic for intrusions and discovery data.

Although an access control policy can inherit its default action from an ancestor policy, you cannot enforce this inheritance.

Security Intelligence

Security Intelligence is a first line of defense against malicious internet content. This feature allows you to block connections based on the latest IP address, URL, and domain name reputation intelligence. To ensure continual access to vital resources, you can override Block list entries with custom Do Not Block list entries.

HTTP Responses

When the system blocks a user's website request, you can either display a generic system-provided response page, or a custom page. You can also display a page that warns users, but also allows them to continue to the originally requested site.

Logging

Settings for access control policy logging allow you to configure default syslog destinations for the current access control policy. The settings are applicable to the access control policy and all the included SSL, prefilter, and intrusion policies unless the syslog destination settings are explicitly overridden with custom settings in included rules and policies.

Advanced Access Control Options

Advanced access control policy settings typically require little or no modification. Often, the default settings are appropriate. Advanced settings you can modify include traffic preprocessing, SSL inspection, identity, and various performance options.

System-Created Access Control Policies

Depending on your devices' initial configurations, system-provided policies can include:

- Default Access Control—Blocks all traffic without further inspection.
- Default Intrusion Prevention—Allows all traffic, but also inspects with the Balanced Security and Connectivity intrusion policy and default intrusion variable set.
- Default Network Discovery—Allows all traffic while inspecting it for discovery data but not intrusions or exploits.

Requirements and Prerequisites for Access Control Policies

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin
- You can define custom user roles to differentiate between the intrusion configuration in access control policy and rules and the rest of the access control policy and rules. Using these permissions, you can separate the responsibilities of your network administration team and your intrusion administration teams. The existing pre-defined user roles that included the Modify Access Control Policy permission support all sub-permissions; you need to create your own custom roles if you want to apply granular permissions. The granular permissions are:
  - **Policies** > **Access Control** > **Access Control Policy** > **Modify Access Control Policy** > **Modify Threat Configuration** allows the selection of intrusion policy, variable set, and file policy in a rule, the configuration of the advanced options for Network Analysis and Intrusion Policies, the configuration of the Security Intelligence policy for the access control policy, and intrusion actions in the policy default action. If a user has this option only, the user can modify no other part of the policy or rule.
  - **Modify Remaining Access Control Policy Configuration** controls the ability to edit all other aspects of the policy.

Managing Access Control Policies

You can edit system-provided access control policies and create custom access control policies.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**.

At the top of the page, there are convenient links to some related features: Object management, Intrusion policies, Network Analysis policies, DNS policies, and policy Import/Export.

Step 2 Manage access control policies:

- Analyze Policy—Select one or more policy and then click **Analyze Policies** to evaluate access control policies for anomalies such as redundant or shadowed rules, and take action to fix discovered anomalies. The analysis job is sent to the cloud and takes time to complete. See Identifying and Fixing Anomalies with Policy Analyzer & Optimizer, on page 1763.

  The Anomalies column shows the results of the analysis. Click the **% Optimizable** link to see the anomalies, or **Re-Analyze** to run the analysis again. The Last Analyzed column shows when Policy Analyzer & Optimizer was last run.

After completing analysis and optimization, you can download the reports by selecting the following options from the **More** (⋮) menu: **Download Last Policy Analysis**, **Remediation History**.

Note
To use the Policy Analysis feature, you must be using a Cloud-Delivered Firewall Management Center or an connected to Cisco Security Cloud Control (Security Cloud Control). If your setup does not meet requirements, the explanatory dialog that opens when you click this button includes an **Integrate** button to help get you started. Policy Analyzer & Optimizer operates in the cloud only.

- Create—Click **New Policy**; see Creating a Basic Access Control Policy, on page 1748.
- Columns—Click the **Show/Hide Columns** icon above the list of rules to select which information to show in the table. Click **Show All/Hide All** to quickly add or remove all listed columns, excepting name and actions. Click **Default** to undo all of your customizations.
- Inheritance—Click **Plus** next to a policy with descendants to expand your view of the policy's hierarchy.
- Edit—Click **Edit** (✎); see Editing an Access Control Policy, on page 1749
- Delete—Click **Delete** (🗑). You must remove any device assignments before deleting a policy.

  To delete more than one policy at a time, select the check boxes for the policies, then select **Delete Policies** above the table.
- Copy—Select **Clone** from the **More** (⋮) menu. Device assignments are not retained in the copy.
- Report—Select **Generate Report** from the **More** (⋮) menu.. The report is generated as a background process. Go to the message/notifications center and look in the Tasks list. When the report is finished, you can download it from the notification.
- View the Audit Log—Click **Go to Audit Log** from the **More** (⋮).
- Lock or unlock a policy—See Locking an Access Control Policy, on page 1751.

Creating a Basic Access Control Policy

When you create a new access control policy, it contains default actions and settings. After creating the policy, you are immediately placed in an edit session so that you can adjust the policy to suit your requirements.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**.

Step 2 Click **New Policy**.

Step 3 Enter a unique **Name** and, optionally, a **Description**.

Step 4 Optionally, choose a base policy.

If an access control policy is enforced on your domain, this step is not optional. You must choose the enforced policy or one of its descendants as the base policy.

If you select a base policy, the base policy defines the default action and you cannot select a new one in this dialog box. Logging for connections handled by the default action depends on the base policy.

Step 5 When you do not select a base policy, specify the initial **Default Action**:

- **Block all traffic** creates a policy with the **Access Control: Block All Traffic** default action.
- **Intrusion Prevention** creates a policy with the **Intrusion Prevention: Balanced Security and Connectivity** default action, associated with the default intrusion variable set.
- **Network Discovery** creates a policy with the **Network Discovery Only** default action.

When you select a default action, logging of connections handled by the default action is initially disabled. You can enable it later when you edit the policy.

Tip
If you want to trust all traffic by default, or if you chose a base policy and do not want to inherit the default action, you can change the default action later.

Step 6 Optionally, choose the devices to assign to the policy. To narrow the devices that appear, enter a search string. The list includes both devices and device templates.

If you want to deploy this policy immediately, you must perform this step.

Step 7 Click **Save**.

The new policy opens for edit. You can add rules to it and make other changes as needed. See Editing an Access Control Policy, on page 1749 .

Editing an Access Control Policy

When you edit an access control policy, you should lock it to ensure that your changes do not get overridden by another person who might edit it simultaneously.

You can only edit access control policies that were created in the current domain. Also, you cannot edit settings that are locked by an ancestor access control policy.

To protect the privacy of your session, a warning appears after 30 minutes of inactivity on the policy editor. After 60 minutes, the system discards your changes.

Note If you do not lock the policy, consider the following: Only one person should edit a policy at a time, using a single browser window. Although more than one user can edit an access control policy at a time, when one user saves changes, all changes made by the other users are immediately deleted and these users are moved out of edit mode and back to the access control policy listing page. These other users must start again.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**.

Step 2 Click **Edit** (✎) next to the access control policy you want to edit.

If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Edit your access control policy.

Tip
You can operate on multiple rules at one time by selecting their checkboxes in the left column, then selecting the action you want to perform from the **Select Action** drop-down list next to the search box. Bulk editing is available for enabling and disabling, copying, cloning, moving, deleting, and editing rules, or viewing hit counts or related events. You can also remove object overlaps in the selected rules.

You can change the following settings or perform these actions:

- Name and Description—Click **Edit** (✎) next to the name, make your changes, and click **Save**.
- Default Action—Choose a value from the **Default Action** drop-down list.
- Default Action Settings—Click **Cog** (⚙), make your changes, and click **OK**. You can configure settings for logging, the location of an external syslog server or SNMP trap server, and the variable set associated with an intrusion prevention default action.
- Associated Policies—To edit or change policies in the packet flow, click the policy type in the packet flow representation below the policy name. You can select the **Prefilter Rules**, **Decryption**, **Security Intelligence**, and **Identity** policies. When necessary, click **Access Control** to return to the access control rules.
- Policy Assignment—To identify the managed devices targeted by this policy, or enforce this policy in a subdomain, click the **Targeted: x devices** link. You can assign the policy to devices or device templates.
- Rules—To manage access control rules, and to inspect and block malicious traffic using intrusion and file policies, click **Add Rule**, or right-click an existing rule and select **Edit** or another appropriate action. The actions are also available from the **More** (⋮) button for each rule. See Create and Edit Access Control Rules, on page 1783.
- Layout—Use the **Grid/Table View** icon above the list of rules to change the layout. Grid view provides color-coded objects in an easy-to-see layout. Table view provides a summary list so that you can see more rules at once. You can freely switch views without impacting the rules.
- Columns (Table view only)—Click the **Show/Hide Columns** icon above the list of rules to select which information to show in the table. Click **Hide Empty Columns** to quickly remove all columns that have no information, that is, you are not using those conditions in any rule. Click **Revert to Default** to undo all of your customizations.
- Analyze rule logic. You can select the following options from the **Analyze** menu to examine the logic of your rules:
  - **Hit Count**—To view statistics on how many connections matched each rule.
  - **Enable/Disable Rule Conflicts**—Select whether you want to see information on whether rules interfere with each other.
  - **Show Rule Conflicts**—See whether you have redundant or shadowed rules. These conflicts could prevent certain rules from ever being matched by connections, meaning either that you need to fix the match criteria, move the rule, or simply delete the rule.

- **Show Warnings**—See whether there are rules with configuration issues that you need to address.

- Additional Settings—To change additional settings for the policy, select one of the following options from the **More** drop-down arrow at the end of the packet flow line.
  - **Advanced Settings**—To set preprocessing, SSL inspection, identity, performance, and other advanced options. See Access Control Policy Advanced Settings, on page 1756.
  - **HTTP Responses**—To specify what the user sees in a browser when the system blocks a website request. See Choosing HTTP Response Pages, on page 1875.
  - **Inheritance Settings**—To change the base access control policy for this policy, and to enforce this policy's settings in its descendant policies. See Choosing a Base Access Control Policy, on page 1753 and Locking Settings in Descendant Access Control Policies, on page 1754 .
  - **Logging**—To set the default logging options for the policy.

Step 4 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Locking an Access Control Policy

You can lock an access control policy to prevent other administrators from editing it. Locking the policy ensures that your changes will not be invalidated if another administrator edits the policy and saves changes before you save your changes. Without locking, if multiple administrators edit the policy simultaneously, the first person who saves changes wins, and all other users have their changes erased.

The lock is for the access control policy itself. The lock does not apply to objects used in the policy. For example, another user can edit a network object that is used in a locked access control policy. Your lock remains in place until you explicitly unlock the policy, so you can log out and come back to your edits later.

When locked, other administrators have read-only access to the policy. However, other administrators can assign a locked policy to a managed device.

Before you begin

Any user role that has permission to modify the access control policy has permission to lock it, and to unlock a policy that was locked by another user.

However, the ability to unlock a policy that was locked by another administrator is controlled by the following permission: **Policies** > **Access Control** > **Access Control Policy** > **Modify Access Control Policy** > **Override Access Control Policy Lock**.

If you are using custom roles, your organization might have limited your unlocking abilities by not assigning this permission. Without this permission, only the administrator who locks a policy can unlock it.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**.

Step 2 Click **Edit** (✎) next to the access control policy you want to lock or unlock.

The **Lock Status** column shows whether a policy is already locked, and if so, who locked it. An empty cell indicates that the policy is not locked.

If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration. Or, it is locked by another user.

Step 3 Click the lock icon next to the policy name to lock or unlock the policy.

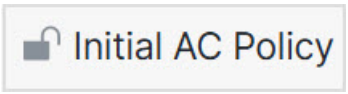

If the policy inherits settings from a parent policy, you must choose one of the following options when you click the lock icon.

- **Lock/Unlock This Policy**—The locking or unlocking is for this policy only.
- **Lock/Unlock This Policy and Parents in the Hierarchy**—This policy and all parent policies are locked or unlocked. If a parent policy is already locked by another administrator, you will see a message and you will not be able to lock that parent policy. When unlocking policies, if you have the Override Access Control Policy Lock permission, all parent policies are unlocked even if they were locked by other users.

Managing Access Control Policy Inheritance

Inheritance relates to using another policy as a base policy for an access control policy. This allows you to use one policy to define some baseline characteristics that can be applied to multiple policies. To understand how inheritance works, see Access Control Policy Inheritance, on page 1731.

Procedure

Step 1 Edit the access control policy whose inheritance settings you want to change; see Editing an Access Control Policy, on page 1749.

Step 2 Manage policy inheritance:

- Change Base Policy — To change the base access control policy for this policy, select **Inheritance Settings** from the **More** drop-down arrow at the end of the packet flow line and proceed as described in Choosing a Base Access Control Policy, on page 1753.
- Lock Settings in Descendants — To enforce this policy's settings in its descendant policies, select **Inheritance Settings** from the **More** drop-down arrow at the end of the packet flow line and proceed as described in Locking Settings in Descendant Access Control Policies, on page 1754 .
- Required in Domains — To enforce this policy in a subdomain, click the **Targeted: x devices** link and proceed as described in Requiring an Access Control Policy in a Domain, on page 1754.

- Inherit Settings from Base Policy — To inherit settings from a base access control policy, click **Security Intelligence**, or select **HTTP Responses** or **Advanced Settings** from the drop-down arrow at the end of the packet flow line, and proceed as directed in Inheriting Access Control Policy Settings from the Base Policy, on page 1753.

Choosing a Base Access Control Policy

You can use one access control policy as the base (parent) for another. By default, a child policy inherits its settings from its base policy, though you can change unlocked settings.

When you change the base policy for the current access control policy, the system updates the current policy with any locked settings from the new base policy.

Procedure

Step 1 In the access control policy editor, select **Inheritance Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 2 Choose a policy from the **Select Base Policy** drop-down list.

Step 3 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Inheriting Access Control Policy Settings from the Base Policy

A new child policy inherits many settings from its base policy. If these settings are unlocked in the base policy, you can override them.

If you later reinherit the settings from the base policy, the system displays the base policy's settings and dims the controls. However, the system saves the overrides you made, and restores them if you disable inheritance again.

Procedure

Step 1 In the access control policy editor, click **Security Intelligence**, or select **HTTP Responses** or **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line

Step 2 Check the **Inherit from base policy** check box for each setting you want to inherit.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration.

Step 3 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Locking Settings in Descendant Access Control Policies

Lock a setting in an access control policy to enforce the setting in all descendant policies. Descendant policies can override unlocked settings.

When you lock settings, the system saves overrides already made in descendant polices so that the overrides can be restored if you unlock settings again.

Procedure

Step 1 In the access control policy editor, select **Inheritance Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 2 In the Child Policy Inheritance Settings area, check the settings you want to lock.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration.

Step 3 Click **OK** to save the inheritance settings.

Step 4 Click **Save** to save the access control policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Requiring an Access Control Policy in a Domain

You can require that every device in a domain use the same base access control policy or one of its descendant policies. This procedure is relevant in a multi-domain deployment only.

Procedure

Step 1 In the access control policy editor, click the **Targeted: x devices** link.

Step 2 Click **Required on Domains**.

Step 3 Build your domain list:

- Add — Select the domains where you want to enforce the current access control policy, then click **Add** or drag and drop into the list of selected domains.
- Delete — Click **Delete** ( ) next to a leaf domain, or right-click an ancestor domain and choose **Delete Selected**.
- Search — Type a search string in the search field. Click **Clear** ( ) to clear the search.

Step 4 Click **OK** to save the domain enforcement settings.

Step 5 Click **Save** to save the access control policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Assigning Devices to an Access Control Policy

An access control policy specifies the devices that use it. Each device can be assigned to only one access control policy. You can also assign the policy to device templates. Templates are included in the list of available and selected devices.

Procedure

Step 1 In the access control policy editor, click the **Targeted: x devices** link.

Step 2 Build your target list:

- Add — Select one or more **Available Devices**, then click **Add to Policy** or drag and drop into the list of **Selected Devices**.
- Delete — Click **Delete** ( ) next to a single device, or select multiple devices, right-click, then choose **Delete Selected**.
- Search — Type a search string in the search field. Click **Clear** ( ) to clear the search.

Under **Impacted Devices**, the system lists the devices whose assigned access control policies are children of the current policy. Any change to the current policy affects these devices.

Step 3 Click **OK** to save your targeted device settings.

Step 4 Click **Save** to save the access control policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Logging Settings for Access Control Policies

To configure logging settings for an access control policy, select **Logging** from the **More** drop-down arrow at the end of the packet flow line.

You can configure default syslog destinations and syslog alert for the access control policy. The settings are applicable to the access control policy and all the included SSL/TLS decryption, prefilter, and intrusion policies unless the syslog destination settings are explicitly overridden with custom settings in included rules and policies.

Logging for connections handled by the default action is initially disabled.

IPS and File and Malware Settings are effective only after you have selected an option at the top of the page for sending syslog messages generally.

Default Syslog Settings

- **Send using specific syslog alert**—If you select this option, the events are sent based on the selected syslog alert as configured using the instructions in *Creating a Syslog Alert Response* in the Cisco Secure Firewall Management Center Administration Guide. You can select the syslog alert from the list or add one by specifying the name, logging host, port, facility, and severity. For more information, see *Facilities and Severities for Intrusion Syslog Alerts* in the Cisco Secure Firewall Management Center Administration Guide. This option is applicable to all devices.

  When using this option, the system sends syslog messages to the server using the Management interface. Ensure there is a route from the Management interface to the syslog server, or messages will not arrive at the server.

- **Use the syslog settings configured in the Threat Defense Platform Settings policy deployed on the device**—If you select this option and select the severity, connection or intrusion events are sent with the selected severity to syslog collectors configured in Platform Settings. Using this option, you can unify the syslog configuration by configuring it in Platform Settings and reusing the settings in access control policy. Severity selected in this section is applied to all connection and intrusion events. The default severity is ALERT.

  This option is applicable only to Secure Firewall Threat Defense devices 6.3 and later.

IPS Settings

- **Send Syslog messages for IPS events**—Send IPS events as syslog messages. The defaults set above are used unless you override them.
- **Show/Hide Overrides**—If you want to use the default syslog destination and severity, leaves these options empty. Otherwise, you can set a different syslog server destination for IPS events, and change the severity of the events.

File and Malware Settings

- **Send Syslog messages for File and Malware events**—Send file and malware events as syslog messages. The defaults set above are used unless you override them.
- **Show/Hide Overrides**—If you want to use the default syslog destination and severity, leaves these options empty. Otherwise, you can set a different syslog server destination for file and malware events, and change the severity of the events.

Access Control Policy Advanced Settings

To configure advanced settings for an access control policy, select **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

Advanced access control policy settings typically require little or no modification. The default settings are appropriate for most deployments. Note that many of the advanced preprocessing and performance options in access control policies may be modified by rule updates as described in *Update Intrusion Rules* in the Cisco Secure Firewall Management Center Administration Guide .

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings.

Caution See Configurations that Restart the Snort Process When Deployed or Activated, on page 256 for a list of advanced setting modifications that restart the Snort process, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

Inheriting Settings from a Parent Policy

If the access control policy has a base policy, you can elect to inherit settings from the base policy. Select **Inherit from base policy** for each setting group where you want to use the parent policy's settings. If inheritance has been configured so that these settings are locked, you cannot configure unique settings for the policy, these settings are read-only.

If you are allowed to configure unique settings for the policy, you must deselect **Inherit from base policy** to make your edits.

General Settings

Option	Description
Maximum URL characters to store in connection events	To customize the number of characters you store for each URL requested by your users. See *Limiting Logging of Long URLs* in the Cisco Secure Firewall Management Center Administration Guide for more information. To customize the length of time before you re-block a website after a user bypasses an initial block, see Setting the User Bypass Timeout for a Blocked Website, on page 1876.
Allow an Interactive Block to bypass blocking for (seconds)	See Setting the User Bypass Timeout for a Blocked Website, on page 1876.
Retry URL cache miss lookup	The first time the system encounters a URL that does not have a locally stored category and reputation, it looks up that URL in the cloud and adds the result to the local data store, for faster processing of that URL in the future. This setting determines what the system does when it needs to look up a URL's category and reputation in the cloud. By default, this setting is enabled: The system momentarily delays the traffic while it checks the cloud for the URL's reputation and category, and uses the cloud verdict to handle the traffic. If you disable this setting: When the system encounters a URL that is not in its local cache, the traffic is immediately passed and handled according to the rules configured for Uncategorized and reputationless traffic. In passive deployments, the system does not retry the lookup, as it cannot hold packets.

Option	Description
Enable Threat Intelligence Director	Disable this option to stop publishing TID data to your configured devices.
Enable reputation enforcement on DNS traffic	This option is enabled by default, for improved URL filtering performance and efficacy. For details and additional instructions, see DNS Filtering: Identify URL Reputation and Category During DNS Lookup , on page 1870 and subtopics.
Inspect traffic during policy apply	To inspect traffic when you deploy configuration changes unless specific configurations require restarting the Snort process, ensure that **Inspect traffic during policy apply** is set to its default value (enabled). When this option is enabled, resource demands could result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Scenarios, on page 252 for more information.

Associated Policies

Use advanced settings to associate subpolicies (decryption, identity, prefilter) with access control; see Associating Other Policies with Access Control, on page 1762.

TLS Server Identity Discovery

The latest version of the Transport Layer Security (TLS) protocol 1.3, defined by RFC 8446, is the preferred protocol for many web servers to provide secure communications. Because the TLS 1.3 protocol encrypts the server's certificate for additional security, and the certificate is needed to match application and URL filtering criteria in access control rules, the Firepower System provides a way to extract the server certificate *without* decrypting the entire packet.

You can enable this feature, referred to as *TLS server identity discovery*, when you configure advanced settings for an access control policy. Certain features are not supported, such as STARTTLS traffic, the HTTP CONNECT method, and in a network where another device is already decrypting traffic.

If you enable this option, we recommend you also enable the decryption policy's advanced TLS adaptive server identity probe option as well. Together, these options enable more efficient decryption of TLS 1.3 traffic. For more information, see TLS 1.3 Decryption Best Practices, on page 2300.

When a new connection starts that will be affected by TLS server identity discovery, the Firewall Threat Defense holds the original ClientHello packet to determine the identity of the server to which it connects before continuing. The Firewall Threat Defense device sends a specialized connection from the Firewall Threat Defense to the server. The server's response includes the server certificate, the specialized connection is terminated, and the original connection is evaluated as required by the access control policy.

TLS server identity discovery prioritizes the certificate's Common Name (CN) over the Server Name Indication (SNI).

To enable TLS server identity discovery, click the **Advanced** tab, click **Edit** (✎) for the setting, and select **Early application detection and URL categorization**.

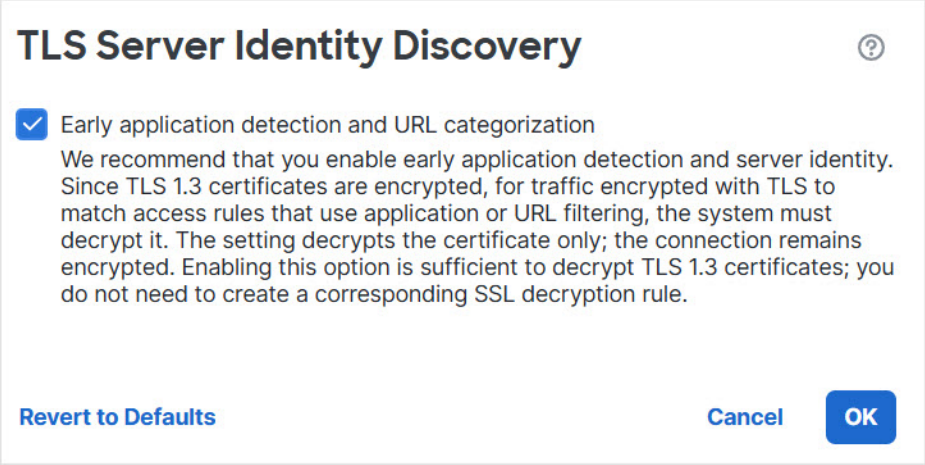

We strongly recommend enabling it for any traffic you want to match on application or URL criteria, especially if you want to perform deep inspection of that traffic. A decryption policy is not required because *traffic is not decrypted* in the process of extracting the server certificate.

Note

- TLS server identity discovery *cannot* be used with any of the following:
  - STARTTLS traffic
  - The HTTP CONNECT method
  - Traffic that is already being decrypted by another device on the network
- Because the certificate is decrypted, TLS server identity discovery can reduce performance depending on the hardware platform.
- TLS server identity discovery is not supported in inline tap mode or passive mode deployments.
- Enabling TLS server identity discovery is not supported on any Secure Firewall Threat Defense Virtual deployed to AWS. If you have any such managed devices managed by the Secure Firewall Management Center, the connection event **PROBE_FLOW_DROP_BYPASS_PROXY** increments every time the device attempts to extract the server certificate.
- TLS Server Identity Discovery also operates on TLS 1.2 sessions.

Network Analysis and Intrusion Policies

Advanced network analysis and intrusion policy settings allow you to:

- Specify the intrusion policy and associated variable set that are used to inspect packets that must pass before the system can determine exactly how to inspect that traffic.
- Change the access control policy's default network analysis policy, which governs many preprocessing options.
- Use custom network analysis rules and network analysis policies to tailor preprocessing options to specific security zones, networks, and VLANs.

For more information, see Advanced Access Control Settings for Network Analysis and Intrusion Policies, on page 2715.

Threat Defense Service Policy

You can use the Threat Defense Service Policy to apply services to specific traffic classes. For example, you can use a service policy to create a timeout configuration that is specific to a particular TCP application, as opposed to one that applies to all TCP applications. This policy applies to Firewall Threat Defense devices only, and will be ignored for any other device type. The service policy rules are applied after the access control rules. For more information, see Service Policies, on page 1937.

File and Malware Settings

Tuning File and Malware Inspection Performance and Storage, on page 2244 provides information on performance options for file control and malware defense.

Portscan Threat Detection

Portscan detector is a threat detection mechanism designed to help you detect and prevent portscan activity in all types of traffic to protect networks from eventual attacks. Portscan traffic can be detected efficiently in both allowed and denied traffic. For more information, see Threat Detection, on page 1955.

Elephant Flow Settings

Elephant flows are large, long duration, and fast flows that can cause duress for Snort cores. There are two actions that can be applied on elephant flows to reduce system stress, CPU hogging, packet drops, and so on. These actions are:

- Bypass any or all applications—This action bypasses flow from Snort inspection.
- Throttle—This action applies dynamic rate limit policy (10% reduction) on elephant flows.

For more information, see the Elephant Flow Detection chapter in the Cisco Secure Firewall Management Center Snort 3 Configuration Guide.

Intelligent Application Bypass Settings

Intelligent Application Bypass (IAB) is an expert-level configuration that specifies applications to bypass or test for bypass if traffic exceeds a combination of inspection performance and flow thresholds. For more information, see Intelligent Application Bypass, on page 1965.

Transport/Network Layer Preprocessor Settings

Advanced transport and network preprocessor settings apply globally to all networks, zones, and VLANs where you deploy your access control policy.

- **Ignore the VLAN header when tracking connections** —Different VLAN tags in traffic traveling in different directions for the same connection can affect traffic reassembly and rule processing. For example, traffic for the same connection could be transmitted over VLAN A and be received over VLAN B. Select this option to configure the system to ignore the VLAN header so packets can be correctly processed for your deployment.
- **Maximum Active Responses**—For a TCP connection that triggers a preprocessor/intrusion rule that is configured to provide an active response, the maximum number of active responses per TCP connection. When additional traffic occurs on a connection where an active response has been initiated, and the traffic

occurs more than **Minimum Response Seconds** after a previous active response, the system sends another active response unless the specified maximum has been reached. A setting of 0 disables additional active responses triggered active response rules.

- **Minimum Response Seconds**—Until **Maximum Active Responses** occur, specifies the number of seconds to wait before any additional traffic on a connection where the system has initiated an active response results in a subsequent active response.
- **Session Termination Logging Threshold**—Do not modify this option unless instructed to do so by Cisco Technical Support. This option specifies for the number of bytes that result in a logged message when the session terminates and the specified number was exceeded. Changing the option can affect system performance.

Detection Enhancement Settings

Detection Enhancement Settings determine whether adaptive profiles are used for application detection and intrusion rules in the access control policy. Typically, the system uses the static settings in your network analysis policy to preprocess and analyze traffic. With adaptive profile updates, the system can adapt processing behavior using host information either detected by network discovery or imported from a third party.

To enable Adaptive Profiles in Snort 3, you must select both the **Enable** and **Enable Profile Updates** options.

- **Enable**—You must enable adaptive profiling (its default state) for access control rules to perform application and file control, including malware protection (AMP), and for intrusion rules to use service metadata.
- **Enable Profile Updates**—Profile updates, like the target-based profiles you can configure manually in a network analysis policy, help to defragment IP packets and reassemble streams in the same way as the operating system on the target host. The intrusion rules engine then analyzes the data in the same format as that used by the destination host. Profile updates also compare metadata in an intrusion rule to host information to determine whether a rule should apply for a particular host.
- **Adaptive Profiles – Attribute Update Interval**—When profile updates are enabled, you can control how frequently in minutes network map data is synced from the management center to its managed devices. The system uses the data to determine what profiles should be used when processing traffic. Increasing the value for this option can improve performance in a large network.
- **Adaptive Profiles – Networks**—Optionally, when profile updates are enabled, you can improve performance by constraining profile updates to a comma-separated list of IP addresses, address blocks, and network variables. If you use a network variable, the system uses the variable's value in the variable set linked to the default intrusion policy for your access control policy. For example, you could enter: 192.168.1.101, 192.168.4.0/24, $HOME_NET. IPv4 and IPv6 are supported.

  The default value (0.0.0.0/0) applies adaptive profile updates to all networks.

Performance Settings and Latency-Based Performance Settings

About Intrusion Prevention Performance Tuning, on page 2099 provides information on improving the performance of your system as it analyzes traffic for attempted intrusions.

For information specific to latency-based performance settings, see Packet and Intrusion Rule Latency Threshold Configuration, on page 2104.

Encrypted Visibility Engine

For details about this feature, see the Encrypted Visibility Engine chapter in the Cisco Secure Firewall Management Center Snort 3 Configuration Guide.

Associating Other Policies with Access Control

The easiest way to associate the major policy's to an access control policy is by clicking the policy's link in the packet flow shown at the topic of the access control policy. You can quickly select the associated policy. Alternatively, you can use the policy's advanced settings to associate the policy, as described in this topic. These policies include the following:

- Prefilter policy—Performs early traffic handling using limited network (layer 4) outer-header criteria.
- Decryption policy—Monitors, decrypts, blocks, or allows application layer protocol traffic encrypted with Secure Socket Layer (SSL) or Transport Layer Security (TLS).

Caution *Snort 2 only*. Adding or removing an SSL policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

- Identity policy—Performs user identification based on the realm and authentication method associated with the traffic.

Before you begin

Before associating an SSL policy with an access control policy, review the information about TLS server identity discovery in Access Control Policy Advanced Settings, on page 1756.

Procedure

Step 1 In the access control policy editor, select **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 2 Click **Edit** ( ) in the appropriate Policy Settings area.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Choose a policy from the drop-down list.

If you choose a user-created policy, you can click edit that appears to edit the policy.

Step 4 Click **OK**.

Step 5 Click **Save** to save the access control policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Identifying and Fixing Anomalies with Policy Analyzer & Optimizer

You can use the Policy Analyzer & Optimizer to evaluate access control policies for anomalies such as redundant or shadowed rules, and take action to fix discovered anomalies. The Policy Analyzer & Optimizer is hosted in the cloud and is different from the rule analysis available when you are not integrated with the cloud. The non-cloud policy analysis is not available once you integrate with the cloud.

The system automatically performs policy analysis on a daily basis (every 24 hours). You can also manually start an analysis. When you initially enable the service, the system starts an analysis of all existing access control policies.

Note Before optimizing a policy, create a copy of the policy. If you are then dissatisfied with the results of optimization, you can easily reassign the managed devices to the copy and return the system to its starting state.

Before you begin

- You can use Policy Analyzer & Optimizer directly from Security Cloud Control with Firewall Management Center 7.2+, but you can cross-launch the feature starting with 7.6 only. The software version running on the managed devices assigned to the access control policies does not matter. Policy Analyzer & Optimizer is hosted in the cloud only, so whether you run the analysis from within the Firewall Management Center or directly from Security Cloud Control does not make a difference.
- To use the feature, you must select **Enable Policy Analysis & Optimization** when you integrate with the Cisco Security cloud, under **Integration** > **Cisco Security Cloud**.
- If you have enabled Change Management, Policy Analyzer & Optimizer automatically creates a ticket for the changes, and submits the ticket. The approver must approve the ticket before the changes can be deployed.
- If you are using domains, you cannot click the link in the Anomaly column to see the report. Instead, log into Security Cloud Control and use the feature from that application.
- Policy Analyzer & Optimizer adds rule comments on rules that are updated, disabled, or merged. You can later search on these comments to find optimized rules.
- Changes implemented by Policy Analyzer & Optimizer are reflected in the audit log as API calls under the default name internaladmin.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**.

If you have already run an analysis, the Anomaly column shows the number of issues with the policy and the percentage the policy can be optimized, and the state of the policy analysis, such as Error or Completed. Last Analyzed shows the date/time when the analysis was run.

Step 2 Select one or more policy, then click **Analyze Policy**.

The analysis runs as a background process in the cloud. When the analysis is complete, the results appear in the Anomaly column.

Notes:

- You can also start an analysis when editing the policy by selecting **Analyze** > **Policy**. Other options from that menu allow you to show hit counts and warnings.
- If you have not connected to the cloud yet, the explanatory dialog that opens when you click this button includes an **Integrate** button to help get you started. Policy Analyzer & Optimizer operates in the cloud only.

Step 3 When the analysis is complete, click the **% Optimizable** link in the Anomaly column to launch Policy Analyzer & Optimizer in the cloud.

When you have done all the actions you want to take, click **Apply Remediations** (in the cloud). You are shown a confirmation of what will be done. Click **Proceed** to implement the changes.

If the initial analysis ended in an error, you could instead click **Re-analyze** to restart the process.

Step 4 Deploy the policy to complete the changes.

If you have Change Management enabled, the approver must first approved the ticket that contains the remediations before you can deploy them.

Viewing Rule Hit Counts

Hit count indicates the number of times a policy rule or default action has been matched to a connection. The hit count is incremented only for the first packet of a connection that matches a rule. You can use this information to identify the efficacy of your rules. Hit count information is available only for access control and prefilter rules applied to Firewall Threat Defense devices.

Note

- The count persists through reboots and upgrades.
- Counts are maintained by each unit in an HA pair or cluster separately.
- You will not be able to derive the hit count information from a device when deployment or a task is in progress on the device.
- You can also see rule hit count information in the device CLI using the **show rule hits** command.
- If you have accessed the Hit Count page from the Access Control Policy page, you will not be able to view or edit prefilter rules and vice-versa.
- Hit counts are not available for rules that use the Monitor action.

Before you begin

If you use custom user roles, ensure that the roles include the following privileges:

- View Device, to see the hit counts.
- Modify Device, to refresh the hit counts.

Procedure

Step 1 In the access control policy or prefilter policy editor , click **Analyze** > **Rule Hit Counts** on the top-right of the page.

Step 2 On the Hit Count page, select the device from the **Select a device** drop-down list.

If it is not the first time that you are generating hit counts for this device, the last fetched hit count information appears next to the drop-down box. Also, verify the **Last Deployed** time to confirm recent policy changes.

Step 3 If necessary, click **Refresh** ( ) to obtain current hit count data from the selected device.

In the prefilter policy, you might need to click **Fetch Current Hit Count** to get initial hit count data.

You cannot refresh the hit count while deployment to the device is in progress.

Step 4 View and analyze the data.

You can do the following:

- Click **Prefilter** or **Access Control** to switch between the hit counts for these policies.
- Search for a specific rule by entering a search string in **Filter** box.
- Broadly limit the list to **Hit Rules** or **Never Hit Rules** by selecting these options in the **Filter by** field. When viewing hit rules, you can further limit the list by selecting a time range in the **In Last** field (for example, in the last 1 day).
- (When viewed from the access control policy.) Clear the hit counts for one or more rule by selecting the checkbox for the rule, then clicking **Clear Hit Counts**. When confirming the action, choose **Clear and Reload** to refresh the hit count data. You can clear hit counts for up to 500 rules at a time. You cannot undo clearing hit counts.

  Note
  Click the checkbox in the table header to select all rules visible in the list. To select a range of rules, select the first rule's checkbox, then Shift+click the last rule's checkbox; all rules in between are also selected.

- (When viewed from the access control policy.) You can do the following with individual rules:
  - Edit the rule by clicking **Edit** from the **More** ( ) menu.
  - Delete the rule from the policy by clicking **Delete** from the **More** ( ) menu.
  - Enable or disable the rule by clicking **Enable/Disable Rule** from the **More** ( ) menu.
  - Clear the hit count (reset it to zero) for the rule by clicking **Clear Hit Count** from the **More** ( ) menu. You cannot undo this action.

- (When viewed from the prefilter policy.) Change the displayed columns by clicking **Cog** ( ) and selecting the columns to show.
- (When viewed from the prefilter policy.) Click on a rule name to edit it, or click **View** ( ) in the last column to view the rule details. Clicking on the rule name highlights it in the policy page where you can edit it.
- (When viewed from the prefilter policy.) Clear the hit count information (reset it to zero) for a rule by right-clicking the rule and selecting **Clear Hit Count**. You can select multiple rules by using Ctrl+click. You cannot undo this action.
- Generate a comma-separated values report of the details on the page by clicking **Generate CSV** on the bottom-left of the page.

Step 5 Click **Close** to return to the policy page.

Analyzing Rule Conflicts and Warnings

Note The feature described in this topic is available only if you do not integrate with the Cisco Security Clound. When you integrate with the cloud, the more powerful Policy Analyzer & Optimizer replaces this feature. See Identifying and Fixing Anomalies with Policy Analyzer & Optimizer, on page 1763.

You can view warnings and information about rule conflicts to examine the logic of your access control policy and to identify rules that need changes. When rules overlap, you can end up with unnecessary rules in the policy, and these rules will never be matched to traffic. The analysis can help you delete unnecessary rules, or identify rules that should be moved or modified so they enforce the desired policy.

Policy warnings and errors point out things that you should understand and perhaps address to ensure that your rules provide the desired services.

Rule conflict analysis identifies the following types of problem:

- Object Overlap—An element in a field of a rule is a subset of one or more elements in the same field of the rule. For example, the source field might include a network object for 10.1.1.0/24, and another object for the host 10.1.1.1. Because 10.1.1.1 is within the network covered by 10.1.1.0/24, the object for 10.1.1.1 is redundant and can be deleted, simplifying the rule and saving device memory.
- Redundant Rule—Two rules apply the same action to the same type of traffic and removing the base rule would not change the ultimate result. For example, if a rule permitting FTP traffic for a particular network were followed by a rule allowing IP traffic for that same network, and there were no rules in between denying access, then the first rule is redundant and you can delete it.
- Shadowed Rule—This is the reverse of a redundant rule. In this case, one rule will match the same traffic as another rule such that the second rule will never be applied to any traffic because it comes later in the access list. If the action for both rules is the same, you can delete the shadowed rule. If the two rules specify different actions for traffic, you might need to move the shadowed rule or edit one of the two rules to implement your desired policy. For example, the base rule might deny IP traffic, and the shadowed rule might permit FTP traffic, for a given source or destination.

Before you begin

When doing the analysis:

- Only the first conflict for a given rule is identified. Once you fix the problem, the rule might be identified as having a conflict with another rule in the table. However, a rule might have multiple warnings or errors.
- Rule conflict analysis considers source/destination security zone, network, VLAN, and service/port match conditions and action only. It does not consider other match criteria, so an apparently redundant rule might not be completely redundant.
- FQDN network objects cannot be analyzed for conflict, because the IP address of an FQDN cannot be known prior to DNS lookup.
- Disabled rules are ignored.
- Time range attributes are ignored. Rules for different time periods might be marked as redundant when they are not actually redundant for the time ranges.
- Icons for warnings and errors, and rule conflicts when you enable the feature, are shown in the rules table. For a reference to the icons, see Rule and Other Policy Warnings, on page 1725.

Procedure

Step 1 Choose **Policy** > **Access Control** and edit an access control policy.

Step 2 Do one of the following to open the rule conflicts and warnings dialog box:

- To view rule conflicts, click the **Analyze** drop-down and click **Enable Rule Conflicts**. When the analysis completes, you see a summary of the conflicts at the top of the page. Then, click **Show Rule Conflicts** from the same menu to see the specific results.

  You must re-enable rule conflict detection each time you open the policy or make a change and save the policy.
- To view rule warnings and errors, click **Analyze** > **Show Rule Warnings**.

  After you make a change to the policy, you can refresh the results by clicking the reload icon next to the Analyze button.
- To view policy warnings, click **Analyze** > **Show Policy Warnings**.
- If you are finished viewing rule conflicts, click **Analyze** > **Disable Rule Conflicts**.

Step 3 In the rule conflicts and warnings dialog box:

- Warnings and Errors are shown on a separate tab from Rule Conflicts.
- Each tab contains sub-tabs to let you examine individual types of problems, such as redundant vs. shadowed or warnings vs. errors. You can also search for an item.
- **More** (⋮) next to each rule name provides shortcuts to edit, disable, or delete the rule.

Step 4 Click **Close** when finished.

Searching for Rules

You can use search to help you find rules, especially when you have a large number of them.

When you search for an IP address in the source or destination network (but not as a simple text search), the system returns rules that will match the address. This includes not only exact matches, but also subnet matches. For example, searching for 10.1.1.1 will include rules for 10.1.1.0/24.

Procedure

Step 1 When editing an access control policy, build the search string by clicking in the **Search** box.

- For a simple text string search, type the string. The search returns rules that have that string in any column. You cannot do a string search in conjunction with a tag search, such as combining a string search with a Source Networks search.
- To search on a specific column, start typing the name of the column until the system prompts you with the full name, such as Source Networks, or select the name from the list of searchable fields. When you select the search tag, you can then enter the search string for that tag. For example, **Source Networks 10.1.1.1**.
- When entering mulitple values, separate the values with commas.
- After your first search, clicking in the search box prompts you with recent searches and tags. You can quickly repeat a search by selecting it, or build similar searches by selecting previous searches or tags and building on them.
- When building a search string with multiple tags, do not include spaces between the tags.
- When you select a tag, you are prompted with values that appear in those columns. Select the values you want to search for.
- You can quickly filter based on some common features by clicking the filter icon to the left of the search box and selecting to show rules with any combination of the following: Allow, Block, Monitor, Intrusion Policy, Time Range, Conflicts, Warnings, Errors, Disabled rules, expired rules, rules that include objects whose definitions overlap..
- To see the rules that apply to a specific device or set of devices, click the filter icon and select the devices. Rules apply to a device if they use security zones that contain at least one interface on the device, or if they do not include security zones.

Step 2 With your cursor at the end of the search string in the search box, press Enter.

The rules that satisfy the search string are highlighted and non-matching rules are hidden. You can deselect **Show Only Matching Rules** to see the entire table, with the highlighted rules within the table. This lets you see the surrounding rules.

Beside the Show Only Matching Rules checkbox is a summary of the total number of rules in the policy compared to the number that match the search string.

Step 3 To close the search and return to the unfiltered and unhighlighted table, click the **X** at the right of the search box. You can also put your cursor at the end of the search string and press the Esc key.

History for Access Control Policies

Table 95:

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Object group search performance enhancements.	7.6.0	7.6.0	Object group search is now faster and uses fewer CPU resources. New CLI commands: **clear asp table network-object**, **show asp table network-object**, **debug acl ogs** Modified CLI comments (enhanced output): , **packet-tracer**, **show access-list**, **show object-group**
Policy Analyzer & Optimizer for access control.	From mgmt. center: 7.6.0 From Security Cloud Control: 7.2.0	Any	The Policy Analyzer & Optimizer evaluates access control policies for anomalies such as redundant or shadowed rules, and can take action to fix discovered anomalies. You can launch the access control Policy Analyzer & Optimizer directly from a Version 7.6+ Firewall Management Center; this requires Cisco Security Cloud. For Versions 7.2–7.4 Firewall Management Centers, use Security Cloud Control. New/modified screens: • To enable: **Integration** > **Cisco Security Cloud** > **Enable Policy Analyzer & Optimizer** • To analyze policies: **Policies** > **Access Control**, select policies, click **Analyze Policies**.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Granular permissions for modifying access control policies and rules.	7.4.0	Any	You can define custom user roles to differentiate between the intrusion configuration in access control policies and rules and the rest of the access control policy and rules. Using these permissions, you can separate the responsibilities of your network administration team and your intrusion administration teams. When defining user roles, you can select the **Policies** > **Access Control** > **Access Control Policy** > **Modify Access Control Policy** > **Modify Threat Configuration** option to allow the selection of intrusion policy, variable set, and file policy in a rule, the configuration of the advanced options for Network Analysis and Intrusion Policies, the configuration of the Security Intelligence policy for the access control policy, and intrusion actions in the policy default action. You can use the **Modify Remaining Access Control Policy Configuration** to control the ability to edit all other aspects of the policy. The existing pre-defined user roles that included the Modify Access Control Policy permission continue to support all sub-permissions; you need to create your own custom roles if you want to apply granular permissions.
New access control policy user interface and rule conflict analysis.	7.3.0	Any	The access control policy user interface introduced in 7.2 is now the default interface. You can also enable rule conflict analysis to help identify redundant rules and objects, and shadowed rules that cannot be matched due to previous rules in the policy.
Access control policy locking.	7.2.0	Any	You can lock an access control policy to prevent other administrators from editing it. Locking the policy ensures that your changes will not be invalidated if another administrator edits the policy and saves changes before you save your changes. Any user who has permission to modify the access control policy has permission to lock it. We added an icon to lock or unlock a policy next to the policy name while editing the policy. In addition, there is a new permission to allow users to unlock policies locked by other administrators: Override Access Control Policy Lock. This permission is enabled by default in the Administrator, Access Admin, and Network Admin roles.
Rule hit counts persist over reboot.	7.2.0	Any	Rebooting a managed device no longer resets access control rule hit counts to zero. Hit counts are reset only if you actively clear the counters. In addition, counts are maintained by each unit in an HA pair or cluster separately. You can use the **show rule hits** command to see cumulative counters across the HA pair or cluster, or see the counts per node. We modified the following device CLI command: **show rule hits**.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Usability improvements for the access control policy.	7.2.0	Any	There is a new user interface available for the access control policy. You can continue to use the legacy user interface, or you can try out the new user interface. The new interface has both a table and a grid view for the rules list, the ability to show or hide columns, enhanced search, infinite scroll, a clearer view of the packet flow related to policies associated with the access control policy, and a simplified add/edit dialog box for creating rules. You can freely switch back and forth between the legacy and new user interfaces while editing an access control policy.
DNS filtering	7.0.0 6.7.0 (experimental)	Any	If URL filtering is enabled and configured, a new option to enhance category and reputation filtering efficacy is enabled by default for each new access control policy. For more information, see DNS Filtering: Identify URL Reputation and Category During DNS Lookup , on page 1870 and subtopics. The Advanced tab of access control policy has a new option under General Settings: **Enable reputation enforcement on DNS traffic**.
TLS server identity discovery	6.7.0	Any	Enable access control policies to evaluate URL and application conditions when a client connects to a TLS 1.3-enabled server. TLS server identity discovery enables these conditions to be evaluated without decrypting traffic. Enabling this feature can impact device performance, depending on model. The Advanced tab page of access control policy has new options: • Warning is displayed on the Advanced tab; moving the slider to the right enables TLS server identity discovery. • New option on the Advanced tab page: **TLS Server Identity Discovery**.
New Security Intelligence categories	—	Any	The following categories were introduced at about the time of the 6.6 release, but are not specific to 6.6: • banking_fraud • high_risk • ioc • link_sharing • malicious • newly_seen • spyware

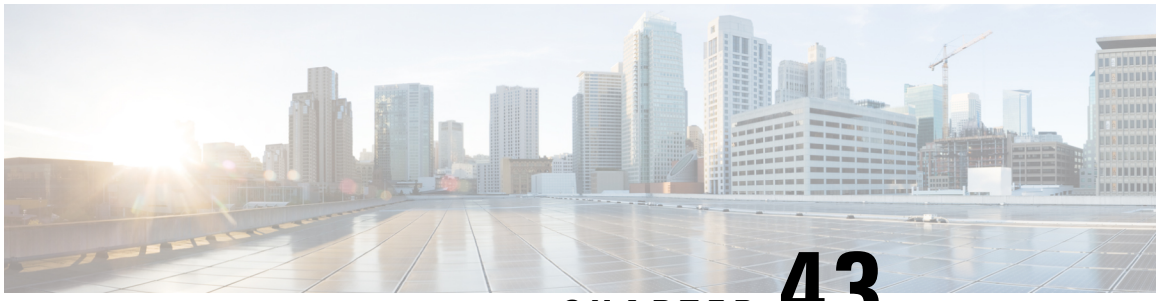

CHAPTER 43

Access Control Rules

The following topics describe how to configure access control rules:

- Introduction to Access Control Rules, on page 1773
- Requirements and Prerequisites for Access Control Rules, on page 1781
- Guidelines and Limitations for Access Control Rules, on page 1782
- Managing Access Control Rules, on page 1783
- Examples for Access Control Rules, on page 1796
- History for Access Control Rules, on page 1801

Introduction to Access Control Rules

Within an access control policy, *access control rules* provide a granular method of handling network traffic across multiple managed devices.

Note Security Intelligence filtering, decryption, user identification, and some decoding and preprocessing occur before access control rules evaluate network traffic.

The system matches traffic to access control rules in the order you specify. In most cases, the system handles network traffic according to the *first* access control rule where *all* the rule's conditions match the traffic.

Each rule also has an *action*, which determines whether you monitor, trust, block, or allow matching traffic. When you allow traffic, you can specify that the system first inspect it with intrusion or file policies to block any exploits, malware, or prohibited files before they reach your assets or exit your network.

The following scenario summarizes the ways that traffic can be evaluated by access control rules in an inline, intrusion prevention deployment.

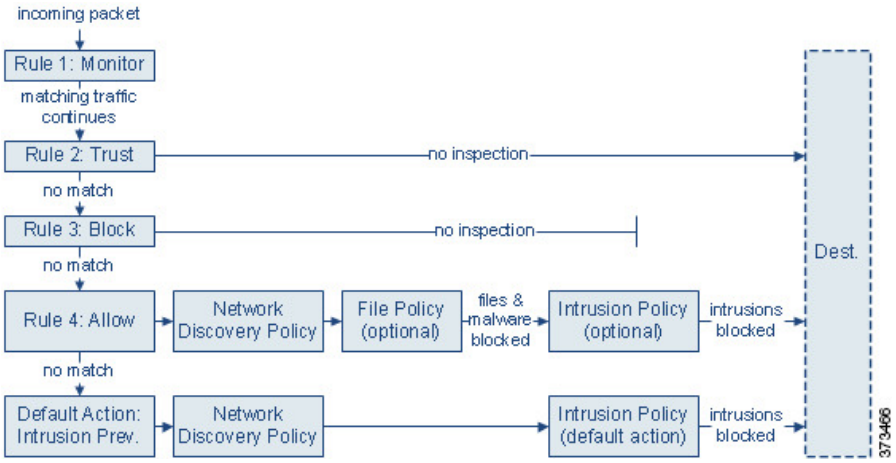

In this scenario, traffic is evaluated as follows:

- **Rule 1: Monitor** evaluates traffic first. Monitor rules track and log network traffic. The system continues to match traffic against additional rules to determine whether to permit or deny it. (However, see an important exception and caveat at Access Control Rule Monitor Action, on page 1778.)
- **Rule 2: Trust** evaluates traffic next. Matching traffic is allowed to pass to its destination without further inspection, though it is still subject to identity requirements and rate limiting. Traffic that does not match continues to the next rule.
- **Rule 3: Block** evaluates traffic third. Matching traffic is blocked without further inspection. Traffic that does not match continues to the final rule.
- **Rule 4: Allow** is the final rule. For this rule, matching traffic is allowed; however, prohibited files, malware, intrusions, and exploits within that traffic are detected and blocked. Remaining non-prohibited, non-malicious traffic is allowed to its destination, though it is still subject to identity requirements and rate limiting. You can configure Allow rules that perform only file inspection, or only intrusion inspection, or neither.
- **Default Action** handles all traffic that does not match any of the rules. In this scenario, the default action performs intrusion prevention before allowing non-malicious traffic to pass. In a different deployment, you might have a default action that trusts or blocks all traffic, without further inspection. (You cannot perform file or malware inspection on traffic handled by the default action.)

Traffic you allow, whether with an access control rule or the default action, is automatically eligible for inspection for host, application, and user data by the network discovery policy. You do not explicitly enable discovery, although you can enhance or disable it. However, allowing traffic does not automatically guarantee discovery data collection. The system performs discovery only for connections involving IP addresses that are explicitly monitored by your network discovery policy; additionally, application discovery is limited for encrypted sessions.

Note that access control rules handle encrypted traffic when your decryption configuration allows it to pass, or if you do not configure decryption. However, some access control rule conditions require unencrypted traffic, so encrypted traffic may match fewer rules. Also, by default, the system disables intrusion and file inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion and file inspection configured.

Access Control Rule Management

The rules table of the access control policy editor allows you to add, edit, categorize, search, filter, move, enable, disable, delete, and otherwise manage access control rules in the current policy.

Properly creating and ordering access control rules is a complex task, but one that is essential to building an effective deployment. If you do not plan your policy carefully, rules can preempt other rules, require additional licenses, or contain invalid configurations. To help ensure that the system handles traffic as you expect, the access control policy interface has a robust warning and error feedback system for rules.

Use the search bar to filter the list of access control policy rules. You can deselect the **Show Only Matching Rules** option to see all rules. Matched rules are highlighted.

For each access control rule, the policy editor displays its name, a summary of its conditions, the rule action, and icons that communicate the rule's inspection options or status. These icons represent:

- **Time Range Option** ()
- **Intrusion policy** ()
- **File policy** ()
- **Logging** ()
- **Warning** ()
- **Errors** ()
- **Rule Conflict** ()

Disabled rules are dimmed and marked (disabled) after the rule name.

To create or edit a rule, use the access control rule editor.

—You can:

- Configure the rule name and select its placement in the upper portion of the editor.
- Switch to editing a different rule by selecting its row above or below the editor.
- Use the left-hand list to select the rule action, and apply intrusion policies and variable sets, file policies, and time range, and to set logging options.
- Use the options next to the rule name to select the rule action, and apply intrusion policies and variable sets, file policies, and time range, and to set logging options.
- Use the **Sources** and **Destinations and Applications** columns to add matching criteria. You can add options from the All list, or move to different tabs to more easily find the type of option you want, such as security zone or networks.
- Add comments to the rule at the bottom of the editor.

Related Topics

Access Control Rule Components, on page 1776
Best Practices for Access Control Rules, on page 1738

Access Control Rule Components

In addition to its unique name, each access control rule has the following basic components:

State

By default, rules are enabled. If you disable a rule, the system does not use it and stops generating warnings and errors for that rule.

Position

Rules in an access control policy are numbered, starting at 1. If you are using policy inheritance, rule 1 is the first rule in the outermost policy. The system matches traffic to rules in top-down order by ascending rule number. With the exception of Monitor rules, the first rule that traffic matches is the rule that handles that traffic.

Rules can also belong to a section and a category, which are organizational only and do not affect rule position. Rule position goes across sections and categories.

Section and Category

To help you organize access control rules, every access control policy has two system-provided rule sections, Mandatory and Default. To further organize access control rules, you can create custom rule categories inside the Mandatory and Default sections.

If you are using policy inheritance, the current policy's rules are nested between its parent policy's Mandatory and Default sections.

Conditions

Conditions specify the specific traffic the rule handles. Conditions can be simple or complex; their use often depends on license.

Traffic must meet all of the conditions specified in a rule. For example, if the Application condition specifies HTTP but not HTTPS, the URL category and reputation conditions will not apply to HTTPS traffic.

Applicable Time

You can specify days and times during which a rule is applicable.

Action

A rule’s action determines how the system handles matching traffic. You can monitor, trust, block, or allow (with or without further inspection) matching traffic. The system does not perform deep inspection on trusted, blocked, or encrypted traffic.

Inspection

Deep inspection options govern how the system inspects and blocks malicious traffic you would otherwise allow. When you allow traffic with a rule, you can specify that the system first inspect it with intrusion or file policies to block any exploits, malware, or prohibited files before they reach your assets or exit your network.

Logging

A rule's logging settings govern the records the system keeps of the traffic it handles. You can keep a record of traffic that matches a rule. In general, you can log sessions at the beginning or end of a connection, or both. You can log connections to the database, as well as to the system log (syslog) or to an SNMP trap server.

Comments

Each time you save changes to an access control rule, you can add comments.

Related Topics

Best Practices for Access Control Rules, on page 1738
Access Control Rule Management, on page 1775
Create and Edit Access Control Rules, on page 1783
Access Control Rule Actions, on page 1778
Access Control Rule Conditions, on page 1785
Deep Inspection Using File and Intrusion Policies, on page 1728
Access Control Rule Comments

Access Control Rule Order

Rules in an access control policy are numbered, starting at 1. The system matches traffic to access control rules in top-down order by ascending rule number.

In most cases, the system handles network traffic according to the *first* access control rule where *all* the rule's conditions match the traffic. Except for Monitor rules, the system does not continue to evaluate traffic against additional, lower-priority rules after that traffic matches a rule.

To help you organize access control rules, every access control policy has two system-provided rule sections, Mandatory and Default. To further organize, you can create custom rule categories inside the Mandatory or Default sections. After you create a category, you cannot move it, although you can delete it, rename it, and move rules into, out of, within, and around it. The system assigns rule numbers across sections and categories.

If you use policy inheritance, the current policy's rules are nested between its parent policy's Mandatory and Default rule sections. Rule 1 is the first rule in the outermost policy, not the current policy, and the system assigns rule numbers across policies, sections, and categories.

Any predefined user role that allows you to modify access control policies also allows you to move and modify access control rules within and among rules categories. You can, however, create custom roles that restrict users from moving and modifying rules. Any user who is allowed to modify access control policies can add rules to custom categories and modify rules in them without restriction.

Caution Failure to set up your access control rules properly can have unexpected results, including traffic being allowed that should be blocked. In general, application control rules should be lower in your access control list because it takes longer for those rules to match than rules based on IP address, for example.

Access control rules that use *specific* conditions (such as networks and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your access control rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your access control rules. For more information about the OSI model, see this Wikipedia article.

Tip Proper access control rule order reduces the resources required to process network traffic, and prevents rule preemption. Although the rules you create are unique to every organization and deployment, there are a few general guidelines to follow when ordering rules that can optimize performance while still addressing your needs.

Related Topics

Best Practices for Ordering Rules, on page 1739

Access Control Rule Actions

Every access control rule has an *action* that determines how the system handles and logs matching traffic. You can monitor, trust, block, or allow (with or without further inspection).

The access control policy's *default action* handles traffic that does not meet the conditions of any access control rule with an action other than Monitor.

Access Control Rule Monitor Action

The **Monitor** action is not designed to permit or deny traffic. Rather, its primary purpose is to force connection logging, regardless of how matching traffic is eventually handled.

If a connection matches a Monitor rule, the next non-Monitor rule that the connection matches should determine traffic handling and any further inspection. If there are no additional matching rules, the system should use the default action.

There is an exception, however. If a Monitor rule contains layer 7 conditions—such as an application condition—the system *allows early packets to pass* and the connection to be established (or the SSL handshake to complete). This occurs even if the connection should be blocked by a subsequent rule; this is because these early packets *are not evaluated against subsequent rules*. So that these packets do not reach their destination completely uninspected, you can specify an intrusion policy for this purpose in the access control policy's Advanced settings; see Inspection of Packets That Pass Before Traffic Is Identified, on page 2716. After the system completes its layer 7 identification, it applies the appropriate action to the remaining session traffic.

Caution As a best practice, *avoid placing layer 7 conditions on broadly-defined monitor rules high in your rule priority order*, to prevent inadvertently allowing traffic into your network. Also, if locally bound traffic matches a Monitor rule in a Layer 3 deployment, that traffic may bypass inspection. To ensure inspection of the traffic, enable **Inspect Local Router Traffic** in the advanced device settings for the managed device routing the traffic.

Access Control Rule Trust Action

The **Trust** action allows traffic to pass without deep inspection or network discovery. Trusted traffic is still subject to identity requirements and rate limiting.

Note

- Some protocols, such as FTP and SIP, use secondary channels, which the system opens through the process of inspection. In some cases, trusted traffic can bypass all inspection, and these secondary channels cannot be opened properly. If you run into this problem, change the trust rule to **Allow**.
- For trust rules with logging options disabled, end-of-flow events are still generated in the system. However, the events are not visible on the event pages.
- Because access control rules are evaluated after other policies, such as decryption, trusting a connection does not necessarily mean it will be fast-pathed with no inspection. For example, if a connection matches both a decryption rule that requires decryption, and a trust access control rule, the connection is decrypted and inspected as needed, prior to being allowed by the trust rule. Trust simply means no additional inspection, such as intrusion inspection, will be applied. If your intention is to allow a connection uninspected, either use the prefilter policy to fast path the connection, or ensure that no other policy applies inspection services to the connection.

Access Control Rule Blocking Actions

The **Block** and **Block with reset** actions deny traffic without further inspection of any kind.

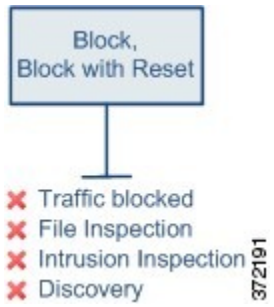

Block with reset rules reset the connection, with the exception of web requests met with an *HTTP response page*. This is because the response page, which you configure to appear when the system blocks web requests, cannot display if the connection is immediately reset.

For more information, see Configure HTTP Response Pages, on page 1873.

Related Topics

Configure HTTP Response Pages, on page 1873

Access Control Rule Interactive Blocking Actions

The **Interactive Block** and **Interactive Block with reset** actions give web users a choice to continue to their intended destinations.

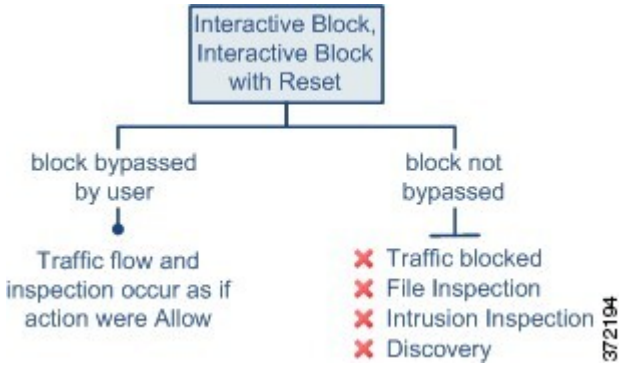

If a user bypasses the block, the rule mimics an allow rule. Therefore, you can associate interactive block rules with file and intrusion policies, and matching traffic is also eligible for network discovery.

If a user does not (or cannot) bypass the block, the rule mimics a block rule. Matching traffic is denied without further inspection.

Note that if you enable interactive blocking, you cannot reset *all* blocked connections. This is because the response page cannot display if the connection is immediately reset. Use the **Interactive Block with reset** action to (non-interactively) block-with-reset all non-web traffic, while still enabling interactive blocking for web requests.

For more information, see Configure HTTP Response Pages, on page 1873.

Related Topics

Decryption Rule Blocking Actions, on page 2336

Access Control Rule Allow Action

The **Allow** action allows matching traffic to pass, though it is still subject to identity requirements and rate limiting.

Optionally, you can use deep inspection to further inspect and block unencrypted or decrypted traffic before it reaches its destination:

- You can use an intrusion policy to analyze network traffic according to intrusion detection and prevention configurations, and drop offending packets depending on the configuration.
- You can perform file control using a file policy. File control allows you to detect and block your users from uploading (sending) or downloading (receiving) files of specific types over specific application protocols.
- You can perform network-based advanced malware protection (AMP), also using a file policy. malware defense can inspect files for malware, and block detected malware depending on the configuration.

The following diagram illustrates the types of inspection performed on traffic that meets the conditions of an Allow rule (or a user-bypassed Interactive Block rule. Notice that file inspection occurs before intrusion inspection; blocked files are not inspected for intrusion-related exploits.

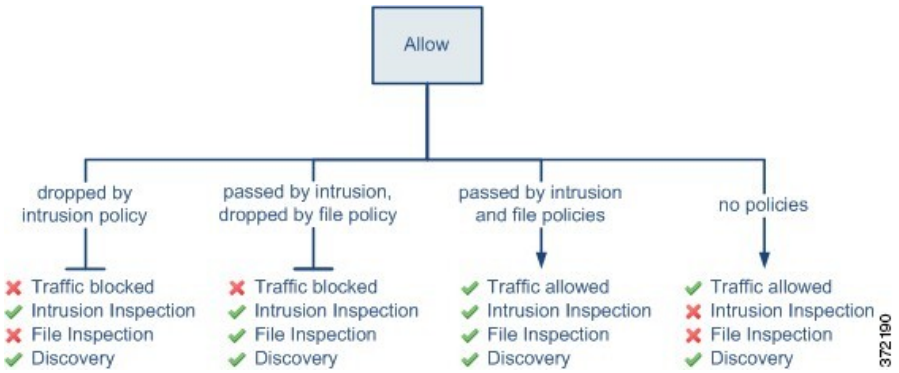

For simplicity, the diagram displays traffic flow for situations where both (or neither) an intrusion and a file policy are associated with an access control rule. You can, however, configure one without the other. Without a file policy, traffic flow is determined by the intrusion policy; without an intrusion policy, traffic flow is determined by the file policy.

Regardless of whether the traffic is inspected or dropped by an intrusion or file policy, the system can inspect it using network discovery. However, allowing traffic does not automatically guarantee discovery inspection. The system performs discovery only for connections involving IP addresses that are explicitly monitored by your network discovery policy; additionally, application discovery is limited for encrypted sessions.

Requirements and Prerequisites for Access Control Rules

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin
- You can define custom user roles to differentiate between the intrusion configuration in access control policy and rules and the rest of the access control policy and rules. Using these permissions, you can separate the responsibilities of your network administration team and your intrusion administration teams. The existing pre-defined user roles that included the Modify Access Control Policy permission support all sub-permissions; you need to create your own custom roles if you want to apply granular permissions. The granular permissions are:
  - **Policies** > **Access Control** > **Access Control Policy** > **Modify Access Control Policy** > **Modify Threat Configuration** allows the selection of intrusion policy, variable set, and file policy in a rule, the configuration of the advanced options for Network Analysis and Intrusion Policies, the configuration of the Security Intelligence policy for the access control policy, and intrusion actions

in the policy default action. If a user has this option only, the user can modify no other part of the policy or rule.

- **Modify Remaining Access Control Policy Configuration** controls the ability to edit all other aspects of the policy.

Guidelines and Limitations for Access Control Rules

- There is a limit of 1000 rules that can be shown at a time on the current page. Thus, if you have a very large number of rules, such as 3000 rules within a single category, actions like selecting all rules in a category and deleting them does not delete all of the rules. You might need to select/delete rules again to delete all the rules that you are targeting.
- If you edit an access control rule that is actively in use, the changes do not apply to established connections at deploy-time. The updated rule is used to match against future connections. However, if the system is actively inspecting a connection (for example, with an intrusion policy), it *will* apply changed matching or action criteria to existing connections.

  For Firewall Threat Defense, you can ensure that your changes apply to all current connections by using the Firewall Threat Defense **clear conn** CLI command to end established connections. Note that you should only do this if it is acceptable to end those connections, on the assumption that the sources for the connections will then attempt to reestablish the connection and thus be matched appropriately against the new rule.
- VLAN tags in access rules only apply to inline sets; they cannot be used in access rules applied to firewall interfaces.
- To use fully-qualified domain name (FQDN) network objects as source or destination criteria, you must also configure DNS for the data interfaces in the platform settings policy. The system does not use the management DNS server setting to do lookups for FQDN objects used in access control rules.

  Note that controlling access by FQDN is a best-effort mechanism. Consider the following points:

  - Whenever possible, use Security Intelligence or URL filtering instead of FQDN rules.
  - Because DNS replies can be spoofed, only use fully trusted internal DNS servers.
  - Some FQDNs, especially for very popular servers, can have hundreds if not thousands of IP addresses, and these can frequently change. Because the system uses cached DNS lookup results, users might get addresses that are not yet in the cache, and their connections will not match the FQDN rule. Rules that use FQDN network objects function effectively only for names that resolve to fewer than 100 addresses.

    We recommend that you do not create network object rules for an FQDN that resolves to more than 100 addresses, as the likelihood of the address in a connection being one that has been resolved and available in the DNS cache on the device is low. For these cases, use a URL-based rule instead of an FQDN network object rule.
  - For popular FQDNs, different DNS servers can return a different set of IP addresses. Thus, if your users use a different DNS server than the one you configure, FQDN-based access control rules might not apply to all IP addresses for the site that are used by your clients, and you will not get the intended results for your rules.

- Some FQDN DNS entries have very short time to live (TTL) values. This can result in frequent recompilation of the lookup table, which can impact overall system performance.
- If more than 8 FQDNs resolve to the same IP address, the system cannot reliably match traffic to the rules for those FQDN. At most 8 FQDNs per IP address can be handled.

- The maximum objects per match criteria per access control rule is 200. For example, you can have up to 200 network objects in a single access control rule.

Managing Access Control Rules

The following topics explain how to manage access control rules.

Adding an Access Control Rule Category

You can divide an access control policy's Mandatory and Default rule sections into custom categories. After you create a category, you cannot move it, although you can delete it, rename it, and move rules into, out of, within, and around it. The system assigns rule numbers across sections and categories.

Procedure

Step 1 In the access control policy editor, click **Add Category**.

Tip
If your policy already contains rules, you can click a blank area in the row for an existing rule to set the position of the new category before you add it. You can also right-click an existing rule and select **Insert new category**.

Step 2 Enter a **Name**.

Step 3 From the **Insert** drop-down list, choose where you want to add the category:

- To insert a category below all existing categories in a section, choose **into Mandatory** or **into Default**.
- To insert a category above an existing category, choose **above category**, then choose a category.
- To insert a category above or below an access control rule, choose **above rule** or **below rule**, then enter an existing rule number.

Step 4 Click **Apply**.

Step 5 Click **Save** to save the policy.

Create and Edit Access Control Rules

Use access control rules to apply actions to specific traffic classes. Rules allow you to selectively allow desirable traffic and drop unwanted traffic.

Procedure

Step 1 In the access control policy editor, you have the following options:

- To add a new rule, click **Add Rule**.
- To edit an existing rule, click **Edit** ( ).
- To start with a copy of an existing rule, select one of the following commands from the **More** ( ) menu.
  - **Copy Rule**, to copy the rule to the clipboard, so you can use the **Paste Above/Below** commands to place it anywhere in the same policy.
  - **Clone Rule**, to create the copy immediately below the duplicated rule.
- To edit multiple rules, use the checkboxes to select multiple rules, then select **Edit** or another action from the **Select Action** list next to the search box.
- To do inline editing, where you change the configuration of an object in a rule condition, right-click the value and choose **Edit**. You can also use the right-click menu to remove a item, add it to the filter, or copy the text or value.

If **View** ( ) appears next to a rule instead, the rule belongs to an ancestor policy, or you do not have permission to modify the rule.

Step 2 If this is a new rule, enter a **Name**.

Step 3 Configure the rule components.

If you are bulk-editing multiple rules, only a subset of options are available.

- Position—Specify the rule position; see Access Control Rule Order, on page 1777.
- Action—Choose a rule **Action**; see Access Control Rule Actions, on page 1778.
- Deep Inspection—(Optional.) For Allow and Interactive Block rules, select options for **Intrusion Policy**, **Variable Set**, and **File Policy**. You can apply intrusion and file policies independently; you do not need to configure both.
- Time Range—(Optional.) For Firewall Threat Defense devices, choose the days and times when the rule is applicable. If you do not choose an option, the rule is always active. For details, see Creating Time Range Objects, on page 1442.
- Logging—Click **Logging** to specify options for connection logging and SNMP traps. See *Best Practices for Connection Logging* in the Cisco Secure Firewall Management Center Administration Guide for more information.
- Conditions—Select the objects you want to add or either source or destination, then click either **Add to Sources** or **Add to Destinations and Applications** to add matching conditions for connections. You can click a tab to restrict the list of available objects, for example, to Networks, Security Zones, Applications, and so forth. However, the sources and destination column always show all selected objects regardless of the tab you are on. See Access Control Rule Conditions, on page 1785 for more information.
- Comments—Open the comments list at the bottom of the dialog box, enter your comment, and click **Post** to add a comment.

Step 4 Click **Add** or **Apply** to save the rule. Click **Apply and Add New Rule** to leave the dialog box open so you can create a new rule.

Step 5 Click **Save** to save the policy.

What to do next

If you will deploy time-based rules, specify the time zone of the device to which the policy is assigned. See Time Zone, on page 1003.

Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Related Topics

Best Practices for Access Control Rules, on page 1738

Access Control Rule Conditions

Rule conditions define the characteristics of the connections you want to target with each rule. Use the conditions precisely to fine-tune the rule to apply to all, and only, the traffic that should be handled by the rule. The following topics explain the match conditions that you can use.

Security/Tunnel Zone Rule Conditions

You can use security zones and tunnel zones to select traffic for a rule.

Security zones segment your network to help you manage and classify traffic flow by grouping interfaces across multiple devices. Tunnel zones allow you to identify tunneled traffic, such as GRE, that should be handled as a tunnel rather than apply access control rules to the encapsulated connections within the tunnel.

You can use security zones to control traffic by its source and destination interfaces. If you add both source and destination zones to a zone condition, matching traffic must originate from an interface in one of the source zones and leave through an interface in one of the destination zones for it to match the rule. Just as all interfaces in a security zone must be of the same type (all inline, passive, switched, or routed), all zones used in a zone condition must be of the same type. Because devices deployed passively do not transmit traffic, you cannot use a zone with passive interfaces as a destination zone.

When using tunnel zones, ensure that you have matching rules in the prefilter policy to associate tunneled traffic with the zone. Then, you can select the tunnel zone as a source zone in a rule; tunnel zones cannot be destinations. If you do not have prefilter rules to rezone the tunnels into the tunnel zone, an access control rule for the tunnel will never apply to any connections. You can specify destination security zones to target tunnels that leave the device through specific interfaces.

Security Zone Considerations

Consider the following when deciding on security zone criteria:

- Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.
- Access control rules generate ACL entries (ACEs) in the device configuration to provide early processing and drops whenever possible. If you specify security zones in rules, ACEs are created for each interface in the zone, which can greatly increase the size of the ACL. Excessively large ACLs generated from access control rules can impact system performance.

Network Rule Conditions

Network rule conditions are the network objects or geographical locations that define the network addresses or locations of the traffic.

- To match traffic from an IP address or geographical location, add the criteria to the Sources list.
- To match traffic to an IP address or geographical location, add the criteria to the Destinations list.
- If you add both source and destination network conditions to a rule, matching traffic must originate from one of the specified IP addresses and be destined for one of the destination IP addresses.

When you add this criteria, you select from the following tabs:

- **Network**—Select the network objects or groups that define the source or destination IP addresses for the traffic you want to control.

  Whenever possible, combine multiple network objects into a single object group. The system automatically creates an object group (during deployment) when you select more than one object (for source or destination separately). Selecting existing groups can avoid object group duplication and reduce the potential impact on CPU usage when there are a large number of duplicate objects.

  You can use objects that define the address using the fully-qualified domain name (FQDN); the address is determined through a DNS lookup. However, FQDN objects are not supported in the following sections in access control policies: Original Client networks, SGT/ISE attributes, Network Analysis And Intrusion policy, Security Intelligence, Threat Detection, Elephant Flow Settings.

- **Geolocation**—Select the geographical location to control traffic based on its source or destination country or continent. Selecting a continent selects all countries within the continent. Besides selecting geographical location directly in the rule, you can also select a geolocation object that you created to define the location. Using geographical location, you could easily restrict access to a particular country without needing to know all of the potential IP addresses used there.

Note To ensure that you are using up-to-date geographical location data to filter your traffic, Cisco strongly recommends that you regularly update the geolocation database (GeoDB).

Original Client in Network Conditions (Filtering Proxied Traffic)

For some rules, you can handle proxied traffic based on the originating client. Use a source network condition to specify proxy servers, then add an original client constraint to specify original client IP addresses. The system uses a packet's X-Forwarded-For (XFF), True-Client-IP, or custom-defined HTTP header field to determine original client IP.

Traffic matches the rule if the proxy's IP address matches the rule's source network constraint, **and** the original client's IP address matches the rule's original client constraint. For example, to allow traffic from a specific original client address, but only if it uses a specific proxy, create three access control rules:

Access Control Rule 1: Blocks proxied traffic from a specific IP address (209.165.201.1)

Source Networks: 209.165.201.1
Original Client Networks: none/any
Action: Block

Access Control Rule 2: Allows proxied traffic from the same IP address, but only if the proxy server for that traffic is one you choose (209.165.200.225 or 209.165.200.238)

Source Networks: 209.165.200.225 and 209.165.200.238
Original Client Networks: 209.165.201.1
Action: Allow

Access Control Rule 3: Blocks proxied traffic from the same IP address if it uses any other proxy server.

Source Networks: any
Original Client Networks: 209.165.201.1
Action: Block

VLAN Tags Rule Conditions

Note VLAN tags in access rules only apply to inline sets. Access rules with VLAN tags do not match traffic on firewall interfaces.

VLAN rule conditions control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic. The system uses the innermost VLAN tag to filter VLAN traffic, with the exception of the prefilter policy, which uses the outermost VLAN tag in its rules.

Note the following Q-in-Q support:

- Firewall Threat Defense on Firepower 4100/9300—Does not support Q-in-Q (supports only one VLAN tag).
- Firewall Threat Defense on all other models:
  - Inline sets and passive interfaces—Supports Q-in-Q, up to 2 VLAN tags.
  - Firewall interfaces—Does not support Q-in-Q (supports only one VLAN tag).

You can use predefined objects to build VLAN conditions, or manually enter any VLAN tag from 1 to 4094. Use a hyphen to specify a range of VLAN tags.

In a cluster, if you encounter problems with VLAN matching, edit the access control policy advanced options, Transport/Network Preprocessor Settings, and select the **Ignore the VLAN header when tracking connections** option.

User Rule Conditions

Matches traffic based the user who initiates the connection, or the group to which the user belongs. For example, you could configure a Block rule to prohibit anyone in the Finance group from accessing a network resource.

You can configure user rule conditions for users in Microsoft Active Directory realms only.

In addition to configuring users and groups for configured realms, you can set policies for the following Special Identities users:

- Failed Authentication: User that failed authentication with the captive portal.
- Guest: Users configured as guest users in the captive portal.

- No Authentication Required: Users that match an identity **No Authentication Required** rule action.
- Unknown: Users that cannot be identified; for example, users that are not downloaded by a configured realm.

For access control rules only, you must first associate an identity policy with the access control policy as discussed in Associating Other Policies with Access Control, on page 1762.

Application Rule Conditions

When the system analyzes IP traffic, it can identify and classify the commonly used applications on your network. This discovery-based *application awareness* is the basis for *application control*—the ability to control application traffic.

System-provided *application filters* help you perform application control by organizing applications according to basic characteristics: type, risk, business relevance, category, and tags. You can create reuseable user-defined filters based on combinations of the system-provided filters, or on custom combinations of applications.

At least one detector must be enabled for each application rule condition in the policy. If no detector is enabled for an application, the system automatically enables all system-provided detectors for the application; if none exist, the system enables the most recently modified user-defined detector for the application. For more information about application detectors, see Application Detector Fundamentals, on page 2624.

You can use both application filters and individually specified applications to ensure complete coverage. However, understand the following note before you order your access control rules.

Benefits of Application Filters

Application filters help you quickly configure application control. For example, you can easily use system-provided filters to create an access control rule that identifies and blocks all high risk, low business relevance applications. If a user attempts to use one of those applications, the system blocks the session.

Using application filters simplifies policy creation and administration. It assures you that the system controls application traffic as expected. Because Cisco frequently updates and adds application detectors via system and vulnerability database (VDB) updates, you can ensure that the system uses up-to-date detectors to monitor application traffic. You can also create your own detectors and assign characteristics to the applications they detect, automatically adding them to existing filters.

Application Characteristics

The system characterizes each application that it detects using the criteria described in the following table. Use these characteristics as application filters.

Table 96: Application Characteristics

Characteristic	Description	Example
Type	Application protocols represent communications between hosts. Clients represent software running on a host. Web applications represent the content or requested URL for HTTP traffic.	HTTP and SSH are application protocols. Web browsers and email clients are clients. MPEG video and Facebook are web applications.
Risk	The likelihood that the application is being used for purposes that might be against your organization's security policy.	Peer-to-peer applications tend to have a very high risk.

Characteristic	Description	Example
Business Relevance	The likelihood that the application is being used within the context of your organization's business operations, as opposed to recreationally.	Gaming applications tend to have a very low business relevance.
Category	A general classification for the application that describes its most essential function. Each application belongs to at least one category.	Facebook is in the social networking category.
Tag	Additional information about the application. Applications can have any number of tags, including none.	Video streaming web applications often are tagged `high bandwidth` and `displays ads`.

Related Topics

Best Practices for Configuring Application Control, on page 1735

Configuring Application Conditions and Filters

To build an application condition or filter, choose the applications whose traffic you want to control from a list of available applications. Optionally (and recommended), constrain the available applications using filters. You can use filters and individually specified applications in the same condition.

Before you begin

- Adaptive profiling must be enabled (its default state) for access control rules to perform application control.

Procedure

Step 1 Invoke the rule or configuration editor:

- Access control, decryption, QoS rule condition—In the rule editor, click **Applications**.
- Identity rule condition—In the rule editor, click **Realms & Settings** and enable active authentication; see Create an Identity Rule, on page 2563.
- Application filter—On the Application Filters page of the object manager, add or edit an application filter. Provide a unique **Name** for the filter.
- Intelligent Application Bypass (IAB)—In the access control policy editor, click **Advanced**, edit IAB settings, then click **Bypassable Applications and Filters**.

Step 2 Find and choose the applications you want to add from the **Available Applications** list.

To constrain the applications displayed in **Available Applications**, choose one or more **Application Filters** or search for individual applications.

Tip

Click **Information**( ) next to an application to display summary information and internet search links. **Unlock** marks applications that the system can identify only in decrypted traffic.

When you choose filters, singly or in combination, the Available Applications list updates to display only the applications that meet your criteria. You can choose system-provided filters in combination, but not user-defined filters.

- Multiple filters for the same characteristic (risk, business relevance, and so on)—Application traffic must match only one of the filters. For example, if you choose both the medium and high-risk filters, the Available Applications list displays all medium and high-risk applications.
- Filters for different application characteristics—Application traffic must match both filter types. For example, if you choose both the high-risk and low business relevance filters, the Available Applications list displays only applications that meet both criteria.

Step 3 Click **Add Application** or **Add to Rule**, or drag and drop.

Tip
Before you add more filters and applications, click **Clear Filters** to clear your current choices.

Step 4 Save or continue editing the rule or configuration.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Port, Protocol, and ICMP Code Rule Conditions

Port conditions match traffic based on the source and destination ports. Depending on the rule type, “port” can represent any of the following:

- TCP and UDP—You can control TCP and UDP traffic based on the port. The system represents this configuration using the protocol number in parentheses, plus an optional associated port or port range. For example: TCP(6)/22.
- ICMP—You can control ICMP and ICMPv6 (IPv6-ICMP) traffic based on its internet layer protocol plus an optional type and code. For example: ICMP(1):3:3.
- Protocol—You can control traffic using other protocols that do not use ports.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Best Practices for Port-Based Rules

Specifying ports is the traditional way to target applications. However, applications can be configured to use unique ports to bypass access control blocks. Thus, whenever possible, use application filtering criteria rather than port criteria to target traffic. Note that application filtering is not available in prefilter rules.

Application filtering is also recommended for applications, like FTP, that open separate channels dynamically for control vs. data flow. Using port-based access control rules can prevent these kinds of applications from performing correctly, and could result in blocking desirable connections.

Using Source and Destination Port Constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as destination port conditions in a single access control rule.

Matching Non-TCP Traffic with Port Conditions

You can match non-port-based protocols. By default, if you do not specify a port condition, you are matching IP traffic. Although you can configure port conditions to match non-TCP traffic, there are some restrictions:

- Access control rules—For Classic devices, you can match GRE-encapsulated traffic with an access control rule by using the GRE (47) protocol as a destination port condition. To a GRE-constrained rule, you can add only network-based conditions: zone, IP address, port, and VLAN tag. Also, the system uses outer headers to match **all** traffic in access control policies with GRE-constrained rules. For Firewall Threat Defense devices, use tunnel rules in the prefilter policy to control GRE-encapsulated traffic.
- Decryption rules—These rules support TCP port conditions only.
- IMCP echo—A destination ICMP port with the type set to 0 or a destination ICMPv6 port with the type set to 129 only matches unsolicited echo replies. ICMP echo replies sent in response to ICMP echo requests are ignored. For a rule to match on any ICMP echo, use ICMP type 8 or ICMPv6 type 128.

URL Rule Conditions

Use URL conditions to control the websites that users on your network can access.

For complete information, see About URL Filtering with Category and Reputation, on page 1858.

Dynamic Attributes Rule Conditions

Dynamic attributes include the following:

- (Source or destination.) Dynamic objects (such as from the Cisco Secure Dynamic Attributes Connector)

  The dynamic attributes connector enables you to collect data (such as networks and IP addresses) from cloud providers and send it to the Secure Firewall Management Center so it can be used in access control rules.

  For more information about the dynamic attributes connector, see About the Cisco Secure Dynamic Attributes Connector, on page 1803.
- (Source only.) Location IP objects, defined by Cisco ISE
- (Source only.) Device type objects, defined by Cisco ISE (also referred to as endpoint profile objects)

Dynamic attributes can be used as source criteria and destination criteria in access control rules. Use the following guidelines:

- Objects of different types are ANDd together
- Objects of a similar type are ORd together

For example, if you choose source destination criteria SGT 1, SGT 2, and device type 1; the rule is matched if device type 1 is detected on either SGT 1 or SGT 2. As another example, if you select both a security group tag, and a dynamic object that lists IP addresses, the rule matches if traffic with the tag originates from (or is destined to) one of those IP addresses.

Time and Day Rule Conditions

You can specify a continuous time range or a recurring time period.

For example, a rule can apply only during weekday working hours, or every weekend, or during a holiday shutdown period.

Time-based rules are applied based on the local time of the device that processes the traffic.

Time-based rules are supported only on Firewall Threat Defense devices. If you assign a policy with a time-based rule to a different type of device, the time restriction associated with the rule is ignored on that device. You will see warnings in this case.

Enabling and Disabling Access Control Rules

When you create an access control rule, it is enabled by default. If you disable a rule, the system does not use it to evaluate network traffic and stops generating warnings and errors for that rule. When viewing the list of rules in an access control policy, disabled rules are grayed out, although you can still modify them.

You can also enable or disable an access control rule using the rule editor.

Procedure

Step 1 In the access control policy editor, right-click the rule and choose a rule state.

If **View** ( ) appears next to a rule instead, the rule belongs to an ancestor policy, or you do not have permission to modify the rule.

Step 2 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Copying Access Control Rules from One Access Control Policy to Another

You can copy access control rules from one access control policy to another. You can copy the rules either to the **Default** section or the **Mandatory** section of the access control policy.

All the settings of the copied rules, except the comments, are retained in the pasted version.

Procedure

Step 1 Do one of the following:

- To copy a single rule, right-click the rule and select **Copy to Different Policy**.
- To copy multiple rules, select their checkboxes, then select **Copy to Different Policy** from the **Select Bulk Action** menu.

Step 2 Select the destination access control policy from the **Access Policy** drop-down list.

Step 3 From the **Place Rules** drop-down list, choose where you want to position the copied rules. You can place them at the top or bottom of either the Mandatory or Default sections.

Step 4 Click **Copy**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Moving Access Control Rules to a Prefilter Policy

You can move access control rules from an access control policy to the associated non-default prefilter policy.

You must first apply a user-defined prefilter policy to the access control policy. The access control rules cannot be moved to the default prefilter policy because the default prefilter policy cannot have rules.

Before you begin

Note the following conditions before you proceed:

- When moving an access control rule to a prefilter policy the layer 7 (L7) parameters in the access control rule cannot be moved. The L7 parameters are dropped during the operation.
- The comments in the access control rule configuration are lost after moving the rule. However, a new comment is added in the moved rule mentioning the source access control policy.
- You cannot move access control rules with **Monitor** set as the **Action** parameter.
- The **Action** parameter in the access control rule is changed to a suitable action in the prefilter rule when moved. To know what each action in the access control rule maps to, see the following table:

Action in the access control rule	Action in the prefilter rule
Allow	Analyze
Block	Block
Block with reset	Block
Interactive Block	Block

Action in the access control rule	Action in the prefilter rule
Interactive Block with reset	Block
Trust	Fastpath

- Similarly, based on the action configured in the access control rule, the logging configuration is set to an appropriate setting after the rule is moved, as mentioned in the following table.

Action in the access control rule	Enabled Logging configurations in the prefilter rule
Allow	None of the check boxes are checked.
Block	• Log at Beginning of Connection • Event Viewer • Syslog Server • SNMP Trap
Block with reset	• Log at Beginning of Connection • Event Viewer • Syslog Server • SNMP Trap
Interactive Block	• Log at Beginning of Connection • Event Viewer • Syslog Server • SNMP Trap
Interactive Block with reset	• Log at Beginning of Connection • Event Viewer • Syslog Server • SNMP Trap
Trust	• Log at Beginning of Connection • Log at End of Connection • Event Viewer • Syslog Server • SNMP Trap

- While moving rules from the source policy, if another user modifies those rules, you will see get a message. You may continue with the process after refreshing the page.

Procedure

Step 1 Do one of the following:

- To move a single rule, right-click the rule and select **Move to Prefilter Policy**.
- To move multiple rules, select their checkboxes, then select **Move to Prefilter Policy** from the **Select Bulk Action** menu.

Step 2 From the **Place Rules** drop-down list, choose where you want to position the moved rules:

- To position as the last set of rules, choose **At the bottom**.
- To position as the first set of rules, choose **At the top**.

Step 3 Click **Move**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Positioning an Access Control Rule

You can move an existing rule within an access control policy, or insert new rules in a desired location. When you add or move a rule to a category, the system places it last in the category.

Before you begin

Review rule order guidelines in Best Practices for Access Control Rules, on page 1738.

Procedure

Step 1 Do one of the following:

- New rule—Insert a new rule by hovering over the line between the existing rules, and clicking **Add Rule**. The location is selected in the **Insert** box in the Add Rule dialog box; you can select a different rule to adjust the location. You can also select **Add Rule Above** or **Add Rule Below** from the right-click menu.
- Existing rules when viewing the rule table—Click and drag the rule to the new poisition.
- Existing rules when viewing the rule table—Right-click a single rule and select **Reposition Rule**. To move multiple rules as a group, select their checkboxes, then select **Reposition Rules** from the **Select Bulk Action** menu.
- Existing rule when editing the rule—Click the **Reposition Rule** icon next to the rule name.

Step 2 Choose where you want to move or insert the rule:

- Choose **into Mandatory** or **into Default**.
- Choose a **into Category**, then choose the category.
- Choose **above rule** or **below rule**, then select the rule.

Step 3 Click **Move** or **Confirm**, and save the rule if you are editing it.

Step 4 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Adding Comments to an Access Control Rule

When you create or edit an access control rule, you can add a comment. For example, you might summarize the overall configuration for the benefit of other users, or note when you change a rule and the reason for the change. You can display a list of all comments for a rule along with the user who added each comment and the date the comment was added.

When you save a rule, all comments made since the last save become read-only.

To search access control rule comments, use the "Search Rules" bar on the rule listing page.

Procedure

Step 1 In the access control rule editor, click **Comments**.

Step 2 Enter your comment and click **Add Comment**. You can edit or delete this comment until you save the rule.

Step 3 Save the rule.

Examples for Access Control Rules

The following topics provide examples of access control rules.

How to Control Access Using Security Zones

Consider a deployment where you want hosts to have unrestricted access to the internet, but you nevertheless want to protect them by inspecting incoming traffic for intrusions and malware.

First, create two security zones: Internal and External. Then, assign interface pairs on one or more devices to those zones, with one interface in each pair in the Internal zone and one in the External zone. Hosts connected to the network on the Internal side represent your protected assets.

Note You are not required to group all internal (or external) interfaces into a single zone. Choose the grouping that makes sense for your deployment and security policies.

Then, configure an access control rule with a destination zone condition set to Internal. This simple rule matches traffic that leaves the device from any interface in the Internal zone. To inspect matching traffic for intrusions and malware, choose a rule action of **Allow**, then associate the rule with an intrusion and a file policy.

How to Control Application Usage

The Web has become the ubiquitous platform for application delivery in the enterprise, whether that is browser based application platforms, or rich media applications that use web protocols as the transport in and out of enterprise networks.

Firewall Threat Defense inspects connections to determine the application being used. This makes it possible to write access control rules targeted at applications, rather than just targeting specific TCP/UDP ports. Thus, you can selectively block or allow web-based applications even though they use the same port.

Although you can select specific applications to allow or block, you can also write rules based on type, category, tag, risk, or business relevance. For example, you could create an access control rule that identifies and blocks all high risk, low business relevance applications. If a user attempts to use one of those applications, the session is blocked.

Cisco frequently updates and adds additional application detectors via system and vulnerability database (VDB) updates. Thus, a rule blocking high risk applications can automatically apply to new applications without you having to update the rule manually.

In this use case, we will block any application that belongs to the **anonymizer/proxy** category.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control** and edit the access control policy.

Step 2 Click **Add Rule** and configure the rule for application control.

a) Give the rule a meaningful name, such as Block_Anonymizers.

b) Select **Block** for **Action**.

c) Assuming that you have configured zones, and want this rule to apply to traffic going from the inside to outside, select the **Zones** tab, and choose your inside zone as the source zone, and the outside zone as the destination zone.

d) Click the **Applications** tab, select the applications to match, and click **Add Application**.

As you select criteria, such as category and risk level, the list to the right of the criteria updates to show exactly which applications match the criteria. The rule you are writing applies to these applications.

Look at this list carefully. For example, you might be tempted to block all very high risk applications. However, as of this writing, TFPT is classified as very high risk. Most organizations would not want to

block this application. Take the time to experiment with various filter criteria to see which applications match your selections. Keep in mind that these lists can change with every VDB update.

For purposes of this example, select anonymizers/proxies from the Categories list and add it to Destinations and Applications. The match criteria should now look like the following graphic.

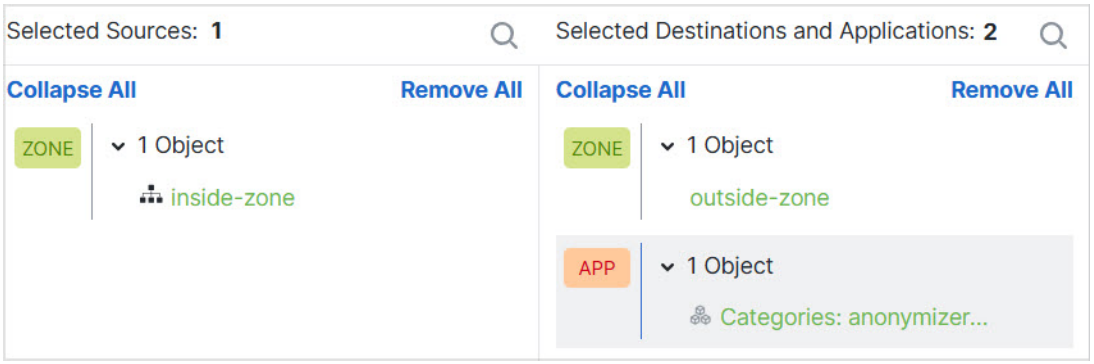

e) Click **Logging** next to the rule action, and enable logging at the start of the connection. You can select a syslog server if you use one.

You must enable logging to get information about any connections blocked by this rule.

Note
To optimize performance, log either the beginning or the end of any connection, but not both. See Cisco Secure Firewall Management Center Administration Guidefor more information.

Step 3 Move the rule so that it comes after any rules that use protocol and port criteria only, but that would not allow traffic that should be blocked by the application rule.

Matching applications requires Snort inspection. Because Snort inspection is not needed by rules that use protocol and port only, you can improve system performance by grouping these simple rules at the top of the access control policy as much as possible.

Step 4 Deploy the changes.

You can use the application rule hit counts and analysis dashboards to see how this rule is performing and how often users try these applications.

How to Block Threats

You can implement next generation Intrusion Prevention System (IPS) filtering by adding intrusion policies to your access control rules. Intrusion policies analyze network traffic, comparing the traffic contents against known threats. If a connection matches a threat you are monitoring, the system drops the connection, thus preventing the attack.

All other traffic handling occurs before network traffic is examined for intrusions. By associating an intrusion policy with an access control rule, you are telling the system that before it passes traffic that matches the access control rule's conditions, you first want to inspect the traffic with an intrusion policy.

You can configure intrusion policies on rules that **allow** traffic only. Inspection is not performed on rules set to **trust** or **block** traffic. In addition, you can configure an intrusion policy as the default action if you do not want to use a simple block.

Besides inspecting traffic that you allow for potential intrusions, you can use the Security Intelligence policy to preemptively block all traffic to or from known bad IP addresses, or to known bad URLs.

This example adds an intrusion policy that allows the internal 192.168.1.0/24 network to got outside, and assumes you already have block rules to selectively eliminate unwanted connections, while also adding a Security Intelligence policy to do pre-emptive blocking.

Before you begin

You must apply the IPS license to any managed device that uses this rule.

This example assumes you have already created security zones for inside and outside interfaces, and the network object for the inside network.

Procedure

Step 1 Create the access control rule that applies the intrusion policy.

a) While editing the access control policy, click **Add Rule**.

b) Give the rule a meaningful name, such as Inside_Outside, and ensure the rule action is **Allow**.

c) For **Intrusion Policy**, select Balanced Security and Connectivity. You can either accept the default variable set or select your own if you want to customize it.

The **Balanced Security and Connectivity** policy is appropriate for most networks. It provides a good intrusion defense without being overly aggressive, which has the potential of dropping traffic that you might not want to be dropped. If you determine that too much traffic is getting dropped, you can ease up on intrusion inspection by selecting the **Connectivity over Security** policy.

If you need to be aggressive about security, try the **Security over Connectivity** policy. The **Maximum Detection** policy offers even more emphasis on network infrastructure security with the potential for even greater operational impact.

If you create your own custom policy, you can select that one instead.

A discussion of variable sets is beyond the scope of this example. Please read the chapters on intrusion policy for detailed information about variable sets and custom policies.

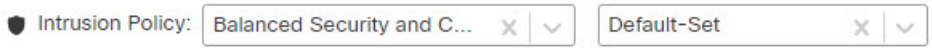

d) Select the **Zones** tab, and add your inside security zone to the source criteria, and outside zone to the destination criteria.

e) Select the **Networks** tab, and add the network object that defines your inside network to the source criteria.

The match criteria should look similar to the following:

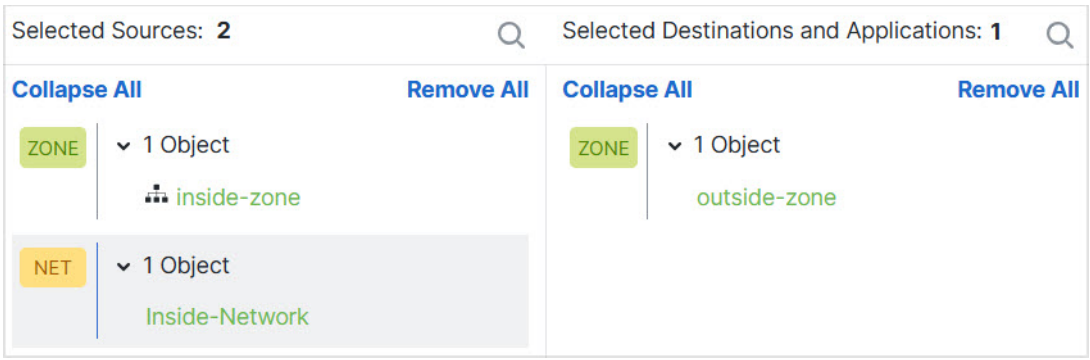

f) Click **Logging** and enable logging at the beginning or end of the connection, or both, as desired.
g) Click **Apply** to save the rule, and then **Save** to save the updated policy.
h) Move the rule to the appropriate location in the access control policy.

Step 2 Configure the Security Intelligence policy to preemptively drop connections with known bad hosts and sites.

By using Security Intelligence to block connections with hosts or sites that are known to be threats, you save your system the time needed to do deep packet inspection to identify threats in each connection. Security Intelligence provides an early block of undesirable traffic, leaving more system time to handle the traffic you really care about.

a) While editing the access control policy, click the **Security Intelligence** link in the packet path.

The link includes two policies: the DNS policy at the top, and the Security Intelligence (Network and URL) at the bottom. In this example, we are configuring the Network and URL lists. By default, these lists already include the global block and do not block lists, but these lists are empty by default until you add items to them.

b) With **Networks** selected and the **Any** security zone selected, scroll down in the list until you get to the global lists, and the first Security Intelligence category (probably Attackers). Click Attackers, then scroll to the end of the categories (probably Tor_exit_node), and Shift+Click to select all of the categories. Click **Add To Block List**.
c) Select the **URL** tab, and **Any** security zone, and use Shift+Click to select URL versions of the same categories. Click **Add To Block List**.
d) Click **Save** to save the policy.
e) As necessary, you can add network and URL objects to the block or do not block lists.

The **Do Not Block** lists are not real "allow" lists. They are exception lists. If an address or URL in the exception list also appears in the blocked list, the connection for the address or URL is allowed to pass on to the access control policy. This way, you can block a feed, but if you later find that a desirable address or site is being blocked, you can use the exception list to override that block without needing to remove the feed entirely. Keep in mind that these connections are subsequently evaluated by access control, and if configured, an intrusion policy. Thus, if any connections do contain threats, they can be identified and blocked during intrusion inspection.

Use the events and dashboards to determine what traffic is actually being dropped by the policy, and whether you need to add addresses or URLs to the **Do Not Block** lists.

Step 3 Deploy your changes.

History for Access Control Rules

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Maximum objects per match criteria per access control rule is 200.	7.3	Any	Previously, you could have a maximum of 50 objects per match criteria in a single access control rule. For example, you could have up to 50 network objects in a single access control rule. The limit is now 200 objects per match criteria in a single rule. We updated the access control policy to allow the increased object limit.
Search for access control rule comments	6.7	Any	The **Search Rules** bar now offers the option to search for comments. New/modified pages: Access control rules page, **Search Rules** text entry field. Supported platforms: Firewall Management Center
Copy or move rules between access control and prefilter policies	6.7	Any	You can copy access control rules from one access control policy to another. You can also move access control rules from an access control policy to the associated prefilter policy. New/modified pages: Access control policy page; the right-click menu for the selected rules provides additional options to copy and move. Supported platforms: Firewall Management Center
Bulk edit of certain settings in access control rules	6.6	Any	In the list of rules in a policy, shift-click or control-click to select multiple rules, then right-click and choose an option. Example bulk operations: You can enable or disable the rules, select a rule action, and edit most inspection and logging settings. New/modified pages: Access control rules page. Supported platforms: Firewall Management Center
Enhanced searching on configured rules	6.6	Any	Enhanced searching on configured rules. New/modified pages: Access control rules page. Supported platforms: Firewall Management Center
Time ranges for rule application	6.6	Any	Ability to specify an absolute or recurring time or time range for a rule to be applied. The rule is applied based on the time zone of the device that processes the traffic. New/modified pages: • A new option on the access control Add Rule page. • A related new option on the **Devices** > **Platform Settings** > **Threat Defense** page to specify time zones for managed devices. Supported platforms: Firewall Threat Defense devices only

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
View object details from access control rule pages	pre-6.6	Any	To see information about an object in a list of rules or from the rule configuration dialog, right-click the object. New/modified pages: **Policies** > **Access Control** > **Access control**, and **Add Rule** page. Supported platforms: Firewall Management Center

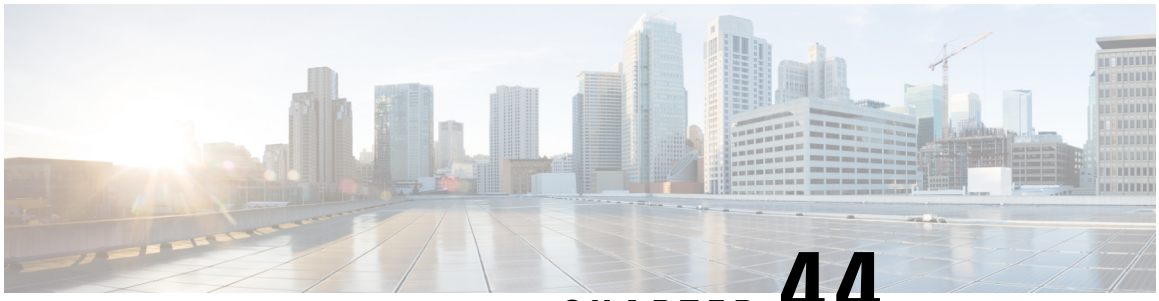

CHAPTER 44

Cisco Secure Dynamic Attributes Connector

The following topics discuss how to configure and use the Cisco Secure Dynamic Attributes Connector.

- About the Cisco Secure Dynamic Attributes Connector, on page 1803
- System Requirements for the Cisco Secure Dynamic Attributes Connector, on page 1806
- Enable the Cisco Secure Dynamic Attributes Connector, on page 1806
- About the Dashboard, on page 1809
- Create a Connector, on page 1816
- Create Dynamic Attributes Filters, on page 1839
- Manually Get a Certificate Authority (CA) Chain, on page 1842
- Use Dynamic Objects in Access Control Policies, on page 1845
- Disable the Cisco Secure Dynamic Attributes Connector, on page 1847
- Troubleshoot Using the Command Line, on page 1848
- Troubleshoot Using the Firewall Management Center, on page 1850
- Manually Get a Certificate Authority (CA) Chain, on page 1850
- Security Requirements, on page 1853
- Internet Access Requirements, on page 1854
- History for the Cisco Secure Dynamic Attributes Connector, on page 1855

About the Cisco Secure Dynamic Attributes Connector

The dynamic attributes connector enables your access control policy to adapt in real time to the changes in public and private cloud workloads and business-critical software-as-a-service (SaaS) applications. It simplifies policy management by keeping rules up to date without tedious manual updates and policy deployment. Customers require policy rules to be defined based on non-network constructs such as VM name or security group, so that firewall policy is persistent even when the IP address or VLAN changes.

Supported connectors

We currently support:

Table 97: List of supported connectors by Cisco Secure Dynamic Attributes Connector version and platform

CSDAC version/platform	AWS	AWS Security Groups	AWS Service Tags	Azure	Azure Service Tags	Cisco APIC	Cisco Cyber Vision	Cisco Multicloud Defense	Generic text	GitHub	Google Cloud	Microsoft Office 365	vCenter	Webex	Zoom
Version 1.1 (on-premises)	Yes	No	No	Yes	Yes	No	No	No	No	No	No	Yes	Yes	No	No
Version 2.0 (on-premises)	Yes	No	No	Yes	Yes	No	No	No	No	No	Yes	Yes	Yes	No	No
Version 2.2 (on-premises)	Yes	No	No	Yes	Yes	No	No	No	No	Yes	Yes	Yes	Yes	No	No
Version 2.3 (on-premises)	Yes	No	No	Yes	Yes	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Version 3.0 (on-premises)	Yes	Yes	Yes	Yes	Yes	No	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Version 3.1 (on-premises)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cloud-delivered (Cisco Security Cloud Control)	Yes	No	No	Yes	Yes	No	No	Yes	No	Yes	Yes	Yes	No	No	No
Secure Firewall Management Center 7.4.1	Yes	No	No	Yes	Yes	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Secure Firewall Management Center 7.6	Yes	Yes	Yes	Yes	Yes	No	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Secure Firewall Management Center 7.7	Yes	Yes	Yes	Yes	Yes	No	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes

How It Works

This topic discusses the architecture of the Cisco Secure Dynamic Attributes Connector.

The following figure shows how the system functions at a high level.

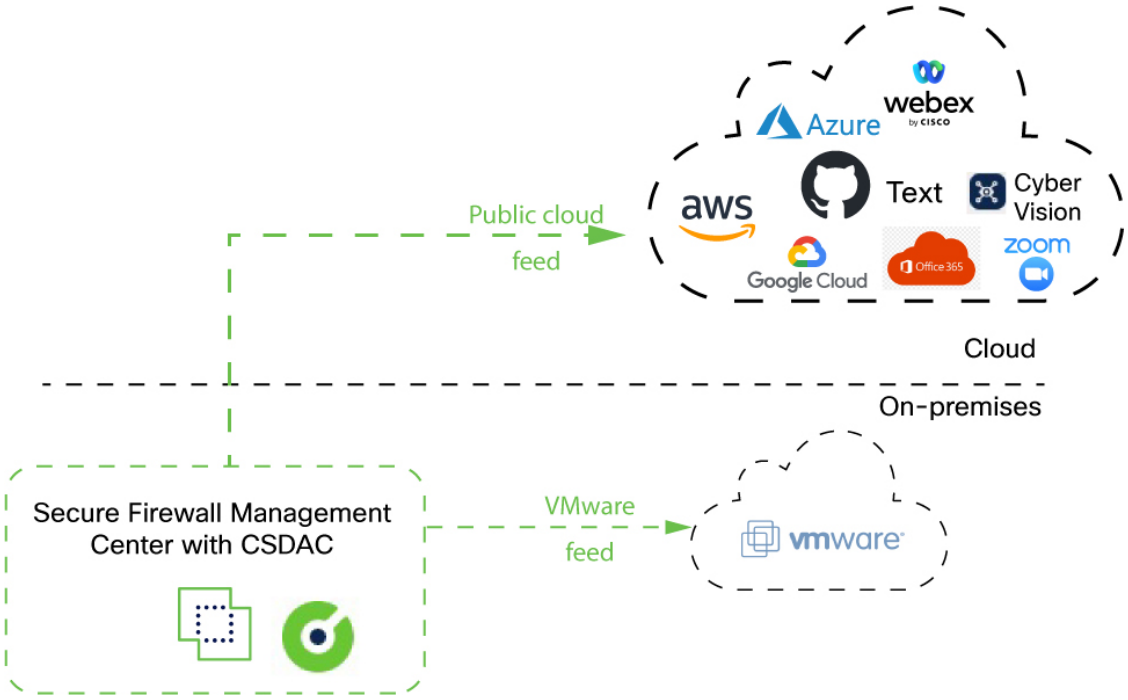

- The system supports certain public cloud providers.

  This topic discusses supported *connectors* (which are the connections to those providers).

- The dynamic attributes connector is provided with Secure Firewall Management Center.

Related topics

- Enable the Cisco Secure Dynamic Attributes Connector, on page 1806
- About the Dashboard, on page 1809

History for the Cisco Secure Dynamic Attributes Connector

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
New connectors	7.6	7.6	AWS security groups, AWS service tags, and Cisco Cyber Vision These connectors can send an on-premises Secure Firewall Management Center dynamic objects as can Cisco Security Cloud Control. To receive dynamic objects from an on-premises dynamic attributes connector, version 3.0 of the on-premises dynamic attributes connector is required.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Cisco Secure Dynamic Attributes Connector	7.4.0	7.4.0	This feature is introduced. The Cisco Secure Dynamic Attributes Connector is now included in the Secure Firewall Management Center. You can use the dynamic attributes connector to get IP addresses from cloud-based platforms such as Microsoft Azure in access control rules without having to deploy to managed devices. More information: • The dynamic attributes connector included with this product: About the Cisco Secure Dynamic Attributes Connector, on page 1803 • The standalone dynamic attributes connector: Cisco Secure Dynamic Attributes Connector Configuration Guide New/modified screen: **Integration** > **Dynamic Attributes Connector**

System Requirements for the Cisco Secure Dynamic Attributes Connector

The Cisco Secure Dynamic Attributes Connector has the following memory requirements:

FMCv: Amount of RAM	Secure Firewall Management Center hardware model	Maximum number of (connectors + Azure AD realms)
At least 32 GB	Firepower 1000, Firepower 1600, vFMC	10
At least 64 GB	Firepower 2500, Firepower 2600,vFMC 300	20
At least 128 GB	Firepower 4500, Firepower 4600	30

The preceding limits apply to both virtual machines and physical machines.

The system prevents you from exceeding these limits to avoid deployment issues.

Enable the Cisco Secure Dynamic Attributes Connector

This task discusses how to enable the Cisco Secure Dynamic Attributes Connector in the Secure Firewall Management Center. The dynamic attributes connector is an integration that enables objects from cloud networking products to be used in Firewall Management Center access control rules.

Procedure

Step 1 Log in to the Secure Firewall Management Center if you have not done so already.

Step 2 Click **Integration** > **Dynamic Attributes Connector**.

Step 3 Slide to **Enabled**.

Step 4 Messages are displayed while the dynamic attributes connector is enabled.

In the event of errors, try again. If errors persist, contact Cisco TAC.

Configure Networks and Subnets for Docker Containers

The Cisco Secure Dynamic Attributes Connector uses Docker containers to retrieve connector data in the Secure Firewall Management Center. To avoid conflicts with the Secure Firewall Management Center management interface and other IP addresses used in your network, you can optionally use the command discussed in this section to change Docker IP addresses and ranges.

About Docker networks

The Docker daemon is used by the dynamic attributes connector requires the following networks:

- `docker0` which is used internally by the Docker daemon.
- A series of IPv6 networks named `veth`*`number`*.

  These are internal bridge networks used by the dynamic attributes connector.
- Docker bridge networks used by dynamic attributes connector connectors named `br-`*`number`*.

Before you enable the dynamic attributes connector, there is only one Docker interface, named `docker0`, set to 172.18.0.1/16 (for a Secure Firewall Management Center Virtual; on-premises managers use different IP address ranges). The table in the Examples section provides details.

Change Docker networks and subnets

First enable the dynamic attributes connector as discussed in Enable the Cisco Secure Dynamic Attributes Connector, on page 1806.

To change Docker networks and subnets, run `/usr/local/sf/bin/change_docker_subnet.sh -b` *`CIDR-network`* `-s` *`address-pool-size`* as a user with `root` privileges where:

- `-b` *`CIDR-network`* sets a network base address pool in CIDR notation.
- `-s` *`address-pool-size`* sets a netmask for the network base address. You can use this option to limit the number of addresses in a base address range in the event the network range overlaps existing network ranges; in particular, we recommend certain `-s` values for Secure Firewall Management Center models to make sure you don't exceed the available RAM in the machine. (Docker containers are used by dynamic attributes connector connectors and those limits are shown in System Requirements for the Cisco Secure Dynamic Attributes Connector, on page 1806.)

Important The networks you assign to Docker must be in an internal network range and must *not* conflict with networks used by the Secure Firewall Management Center or by other devices in your internal network.

Examples

The following table shows examples.

Secure Firewall Management Center model	Recommended -s value	Sample -b value	Cisco Secure Dynamic Attributes Connector container addresses used
Firepower 1000, Firepower 1600, vFMC	27 (netmask 255.255.255.224)	172.19.0.0/16	30 IP addresses `docker0`: 172.19.0.1 Bridge networks `br-number` gateway `172.19.0.33` with subnet `172.19.0.32/27` Connectors created in networks like `172.19.0.38/27`, `172.19.0.39/27`, and so on
Firepower 2500, Firepower 2600, vFMC 300	26 (netmask 255.255.255.192)	192.168.0.0/16	62 IP addresses `docker0`: 192.168.1.1 Bridge networks `br-number` gateway `192.168.1.65`with subnet `192.168.1.64/26` Connectors created in networks like `192.168.1.71/26`, `192.168.1.72/26`, and so on
Firepower 4500, Firepower 4600	25 (netmask 255.255.255.128)	192.168.0.0/16	126 IP addresses `docker0`: 192.168.1.1 Bridge networks `br-number` gateway `192.168.1.129`with subnet `192.168.1.128/25` Connectors created in networks like `192.168.1.136/25`, `192.168.1.135/25`, and so on

For reference, the complete commands follow:

```
sudo /usr/local/sf/bin/change_docker_subnet.sh -b 172.19.0.0/16 -s 27
sudo /usr/local/sf/bin/change_docker_subnet.sh -b 192.168.0.0/16 -s 26
sudo /usr/local/sf/bin/change_docker_subnet.sh -b 192.168.0.0/16 -s 25
```

Verify the networks

To verify your network settings, enter `sudo docker network inspect muster-net`. The command results are displayed in JSON format.

Troubleshoot

Following are some solutions to common errors you might encounter using this command.

Error: `Pull subnet value can not be greater than size`
Solution: Change the value of `-s` so it is less than the CIDR network value.

For example,

INCORRECT: `sudo /usr/local/sf/bin/change_docker_subnet.sh -b 172.19.0.0/16 -s 8`

CORRECT: `sudo /usr/local/sf/bin/change_docker_subnet.sh -b 172.19.0.0/16 -s 20`

Error: After running the command, the Docker networks are wrong.
Solution: Restart the Docker daemon: `sudo pmtool restartbyid docker`

Error: `Cannot connect to the Docker daemon at unix:///var/run/docker.sock. Is the docker daemon running?`
Solution: Restart Docker: `pmtool restartbyid docker`

Error: Input can't be empty
The `-s` parameter is required.

Error: `Pull size - 32 - can not be greater than 32 or less than 0`
Solution: Change the value of `-s` so it is greater than 0 and less than 32.

About the Dashboard

To access the Cisco Secure Dynamic Attributes Connector Dashboard, log in to the Secure Firewall Manager and click **Integration** > **Dynamic Attributes Connector** at the top of the page.

If the Cisco Secure Dynamic Attributes Connector is not enabled, move the slider to enable it. This process could take several minutes to complete.

The Cisco Secure Dynamic Attributes Connector Dashboard page displays the status of your connectors, adapters, and filters at a glance. Following is an example of the Dashboard of an unconfigured system:

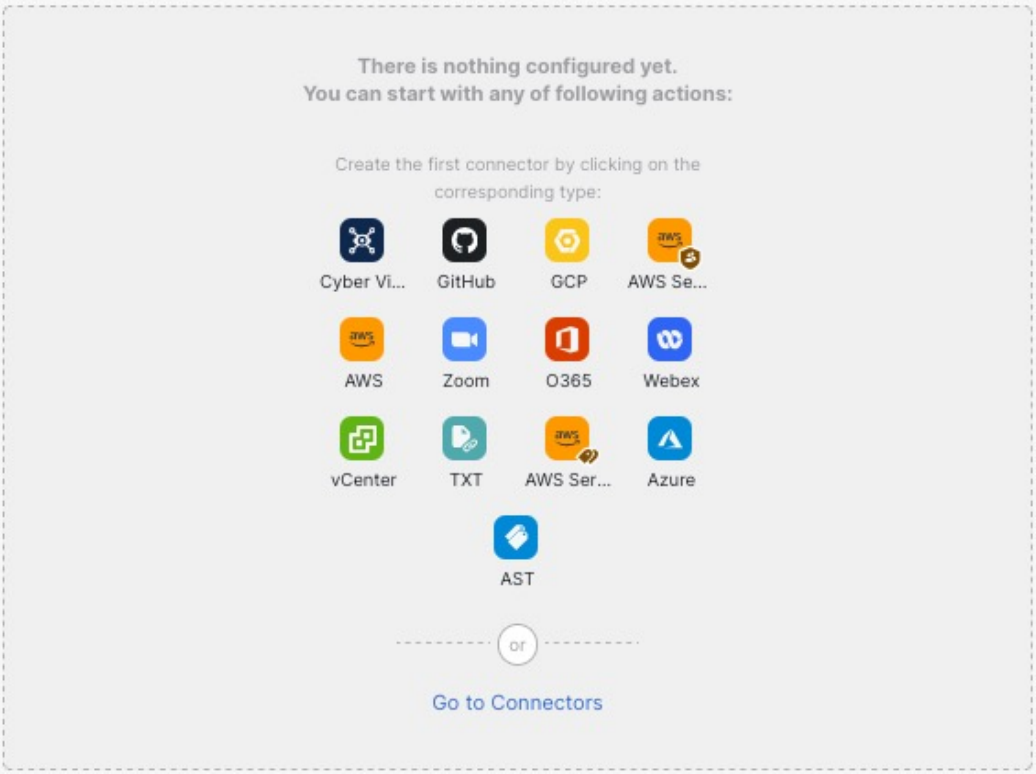

Among the things you can do with the Dashboard are:

- Add, edit, and delete connectors and dynamic attributes filters.
- See how connectors and dynamic attributes filters are related to each other.
- View warnings and errors.

Related Topics

- Dashboard of an Unconfigured System, on page 1810
- Dashboard of a Configured System, on page 1811
- Add, Edit, or Delete Connectors, on page 1813
- Add, Edit, or Delete Dynamic Attributes Filters, on page 1814

Dashboard of an Unconfigured System

Sample Cisco Secure Dynamic Attributes Connector Dashboard page of an unconfigured system:

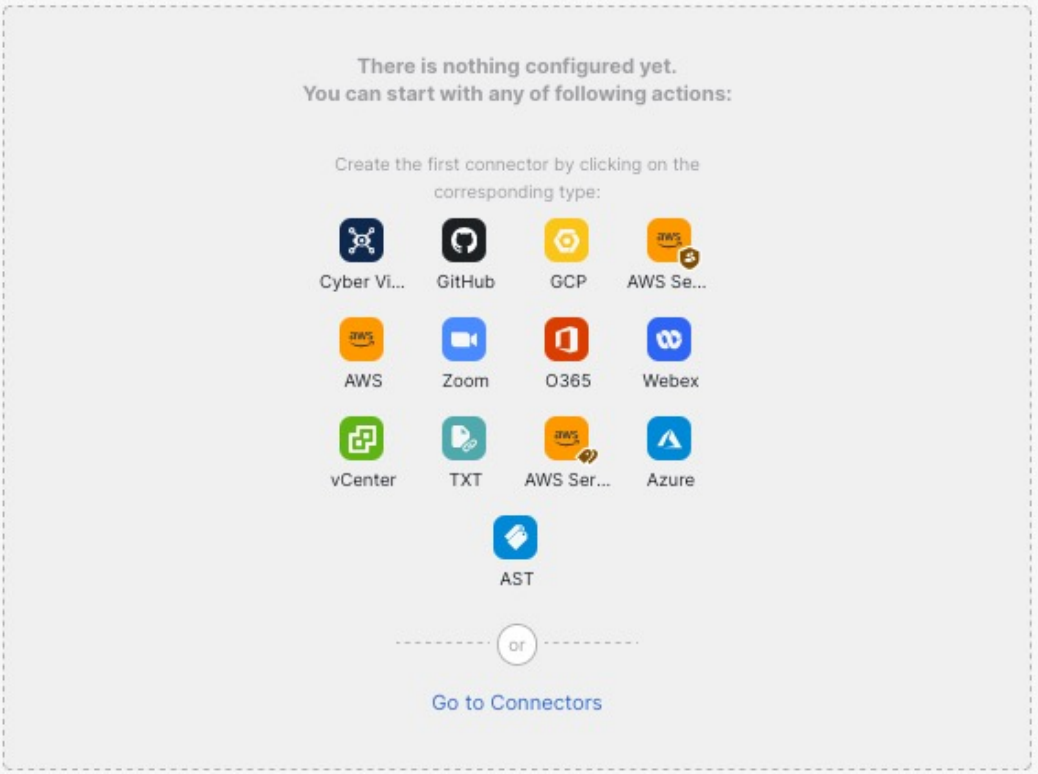

The Dashboard initially displays all the types of connectors you can configure for your system. You can do any of the following:

- Hover the mouse pointer over a connector and click 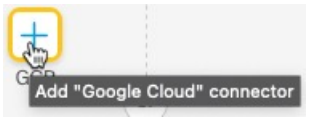

to create a new one.
- Click **Go to Connectors** to add, edit, or delete connectors (good for creating, editing, or deleting multiple connectors at the same time).

  For more information, see Create a Connector, on page 1816.

Related Topics:

- Dashboard of a Configured System, on page 1811
- Add, Edit, or Delete Connectors, on page 1813
- Add, Edit, or Delete Dynamic Attributes Filters, on page 1814

Dashboard of a Configured System

Sample Cisco Secure Dynamic Attributes Connector Dashboard page of a configured system:

Click an area in the figure to learn more about it or click one of the links following the figure.

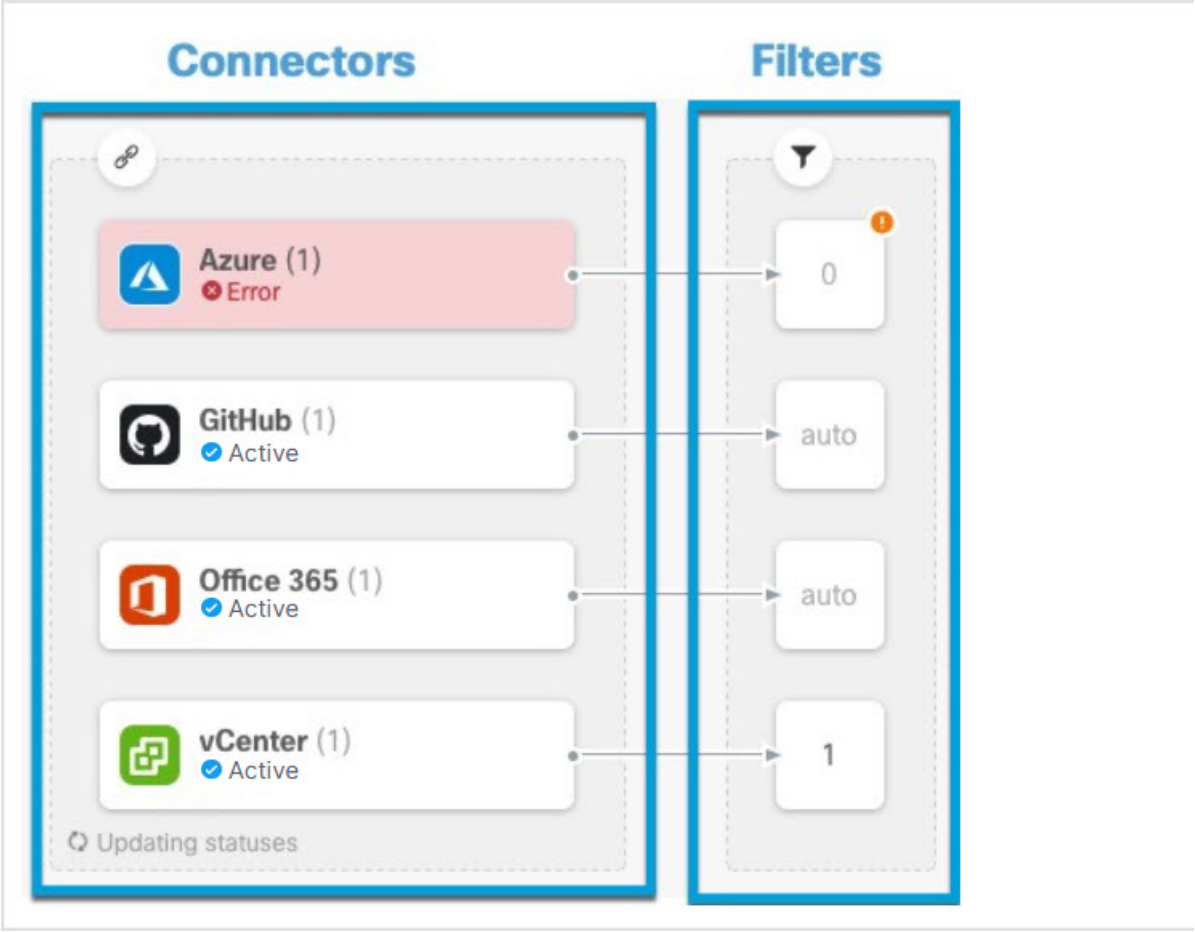

1 Create a Connector, on page 1816
2 Create Dynamic Attributes Filters, on page 1839

The Dashboard shows the following (from left to right):

Connectors column	Filters column
List of connectors with a number indicating how many of each type are configured. Connectors collect dynamic attributes that could be sent to the Secure Firewall Manager. Dynamic attributes filters specify what data is sent. Click to view more information about all configured connectors. You can also click the name of a connector to add, edit, or delete connectors; or to view detailed information about them. For more information, see Add, Edit, or Delete Connectors, on page 1813.	List of dynamic attributes filters associated with each connector with a number indicating how many of each filter are associated with a connector. Click to view more information about all configured filters. You can also click the name of a filter to add, edit, or delete filters; or to view detailed information about them. For more information, see Add, Edit, or Delete Dynamic Attributes Filters, on page 1814.

Note Some connectors, such as Outlook 365 and Azure Service tags, automatically pull available dynamic objects without the need for a dynamic attributes filters. Those connectors display **Auto** in the column.

The Dashboard indicates whether or not an object is available. The Dashboard page is refreshed every 15 seconds but you can click **Refresh** ( ) at the top of the page at any time to refresh immediately. If issues persist, check your network connection.

Related Topics:

- Add, Edit, or Delete Connectors, on page 1813
- Add, Edit, or Delete Dynamic Attributes Filters, on page 1814

Add, Edit, or Delete Connectors

The Dashboard enables you to view or edit connectors. You can click the name of a connector to view all instances of that connector or you can click for the following additional options:

- **Go to Connectors** to view all connectors at the same time; you can add, edit, and delete connectors from there.
- **Add Connector** > ***type*** to add a connector of the indicated type.

Click any connector in the connectors column ( ) to display more information about it; an example follows:

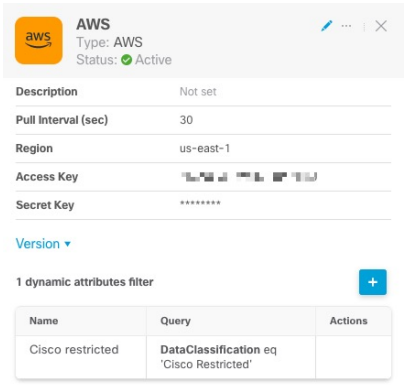

You have the following options:

- Click the Edit icon ( ) to edit this connector.
- Click the More icon ( ) for additional options.
- Click to close the panel.

- Click **Version** to display the version of the . You can optionally copy the version to the clipboard if necessary for Cisco TAC.

The table at the bottom of the panel enables you to add dynamic attributes filters; or to edit or dynamic attributes connector delete connectors. A sample follows:

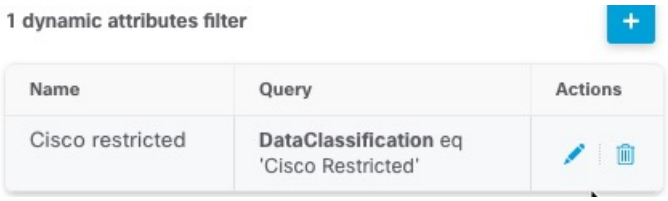

Click the Add icon ( ) to add a dynamic attributes filter for this connector. For more information, see Create Dynamic Attributes Filters, on page 1839.

Hover the mouse pointer over the Actions column to either edit or delete the indicated connector.

View error information

To view error information for a connector:

1. On the Dashboard, click the name of the connector that is displaying the error.
2. In the right pane, click **Information** ( ).

   An example follows.

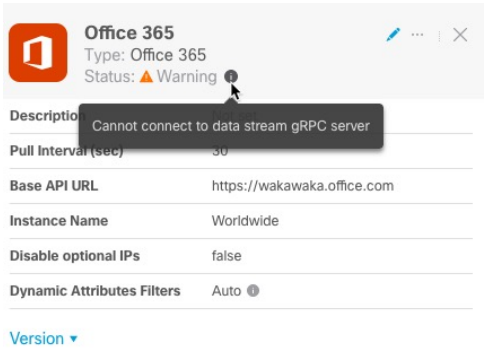

3. To resolve this issue, edit the connector settings as discussed in Create an Office 365 Connector, on page 1833.
4. If you cannot resolve the issue, click **Version** and copy the version to a text file.
5. Provide all of this information to Cisco TAC.

Add, Edit, or Delete Dynamic Attributes Filters

The Dashboard enables you to add, edit, or delete dynamic attributes filters. You can click the name of a filter to view all instances of that filter or you can click for the following additional options:

- **Go to Dynamic Attributes Filters** to view all configured dynamic attributes filters. You can add, edit, or delete dynamic attributes filters from there.
- **Add Dynamic Attributes Filters** to add a filter.

For more information about adding dynamic attributes filters, see Create Dynamic Attributes Filters, on page 1839.

An example follows:

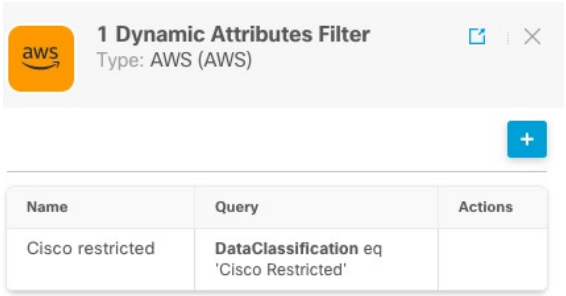

Note Some connectors, such as Outlook 365 and Azure Service tags, automatically pull available dynamic objects without the need for a dynamic attributes filters. Those connectors display **Auto** in the column.

You have the following options:

- Click a filter instance to view summary information about dynamic attributes filters associated with a connector.
- Click the Add icon ( ) to add a new dynamic attributes filter.

  For more information, see Create Dynamic Attributes Filters, on page 1839.
- Click in the filters column ( ) indicates the indicated connector has no associated dynamic attributes filters. Without associated filters, the connector can send nothing to Firewall Management Center.

  One way to resolve the issue is to click in the filters column and click **Add Dynamic Attributes Filter**. A sample follows.

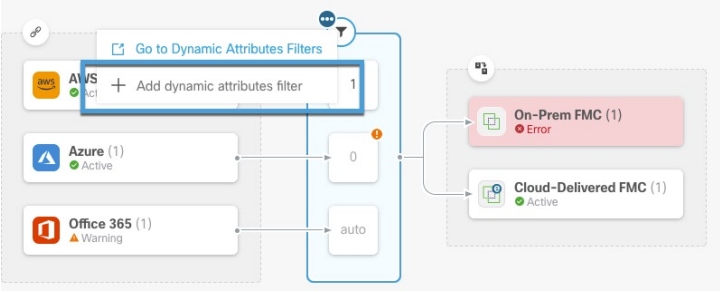

- Click to add, edit, or delete filters.

- Click ✕ to close the panel.

Create a Connector

A *connector* is an interface with a cloud service. The connector retrieves network information from the cloud service so the network information can be used in policies on the Secure Firewall Management Center.

We support the following:

Table 98: List of supported connectors by Cisco Secure Dynamic Attributes Connector version and platform

CSDAC version/platform	AWS	AWS Security Groups	AWS Service Tags	Azure	Azure Service Tags	Cisco APIC	Cisco Cyber Vision	Cisco Multicloud Defense	Generic text	GitHub	Google Cloud	Microsoft Office 365	vCenter	Webex	Zoom
Version 1.1 (on-premises)	Yes	No	No	Yes	Yes	No	No	No	No	No	No	Yes	Yes	No	No
Version 2.0 (on-premises)	Yes	No	No	Yes	Yes	No	No	No	No	No	Yes	Yes	Yes	No	No
Version 2.2 (on-premises)	Yes	No	No	Yes	Yes	No	No	No	No	Yes	Yes	Yes	Yes	No	No
Version 2.3 (on-premises)	Yes	No	No	Yes	Yes	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Version 3.0 (on-premises)	Yes	Yes	Yes	Yes	Yes	No	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Version 3.1 (on-premises)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cloud-delivered (Cisco Security Cloud Control)	Yes	No	No	Yes	Yes	No	No	Yes	No	Yes	Yes	Yes	No	No	No
Secure Firewall Management Center 7.4.1	Yes	No	No	Yes	Yes	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Secure Firewall Management Center 7.6	Yes	Yes	Yes	Yes	Yes	No	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Secure Firewall Management Center 7.7	Yes	Yes	Yes	Yes	Yes	No	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Amazon Web Services Connector—About User Permissions and Imported Data

The Cisco Secure Dynamic Attributes Connector imports dynamic attributes from AWS to Secure Firewall Management Center for use in policies.

Dynamic attributes imported

We import the following dynamic attributes from AWS:

- *Tags*, user-defined key-value pairs you can use to organize your AWS EC2 resources.

  For more information, see Tag your EC2 Resources in the AWS documentation

- *IP addresses* of virtual machines in AWS.

Minimum permissions required

The Cisco Secure Dynamic Attributes Connector requires a user at minimum with a policy that permits `ec2:DescribeTags`, `ec2:DescribeVpcs`, and `ec2:DescribeInstances` to be able to import dynamic attributes.

Create an AWS User with Minimal Permissions for the Cisco Secure Dynamic Attributes Connector

This task discusses how to set up a service account with minimum permissions to send dynamic attributes to Secure Firewall Management Center . For a list of these attributes, see Amazon Web Services Connector—About User Permissions and Imported Data, on page 1816.

Before you begin

You must already have set up your Amazon Web Services (AWS) account. For more information about doing that, see this article in the AWS documentation.

Procedure

Step 1 Log in to the AWS console as a user with the admin role.

Step 2 From the Dashboard, click **Security, Identity & Compliance** > **IAM**.

Step 3 Click **Access Management** > **Users**.

Step 4 Click **Add Users**.

Step 5 In the **User Name** field, enter a name to identify the user.

Step 6 Click **Access Key - Programmatic Access**.

Step 7 At the Set permissions page, click **Next** without granting the user access to anything. You can grant user access later.

Step 8 Add tags to the user if desired.

Step 9 Click **Create User**.

Step 10 Click **Download .csv** to download the user's key to your computer.

Note
This is the only opportunity you have to retrieve the user's key.

Step 11 Click **Close**.

Step 12 At the Identity and Access Management (IAM) page in the left column, click **Access Management** > **Policies**.

Step 13 Click **Create Policy**.

Step 14 On the Create Policy page, click **JSON**.

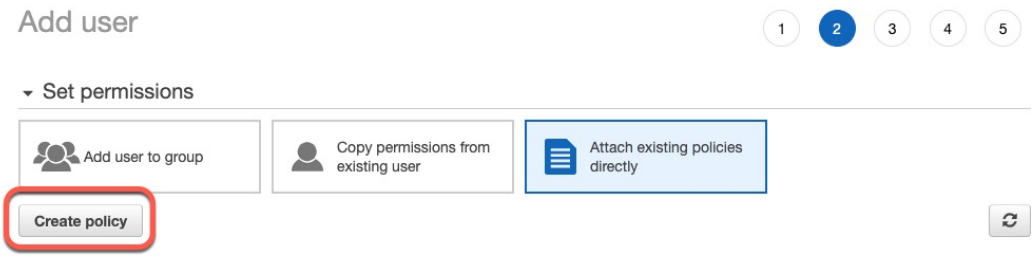

Step 15 Enter the following policy in the field:

```
{
 "Version": "2012-10-17",
 "Statement": [
  {
   "Effect": "Allow",
   "Action": [
    "ec2:DescribeTags",
    "ec2:DescribeInstances",
    "ec2:DescribeVpcs"
   ],
   "Resource": "*"
  }
 ]
}
```

Step 16 Click **Next**.

Step 17 Click **Review**.

Step 18 On the Review Policy page, enter the requested information and click **Create Policy**.

Step 19 On the Policies page, enter all or part of the policy name in the search field and press Enter.

Step 20 Click the policy you just created.

Step 21 Click **Actions** > **Attach**.

Step 22 If necessary, enter all or part of the user name in the search field and press Enter.

Step 23 Click **Attach Policy**.

What to do next

Create an AWS Connector, on page 1818.

Create an AWS Connector

This task discusses how to configure a connector that sends data from AWS to the Secure Firewall Management Center for use in policies.

Before you begin

Create a user with at least the privileges discussed in Create an AWS User with Minimal Permissions for the Cisco Secure Dynamic Attributes Connector, on page 1817.

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Integration** > **Dynamic Attributes Connector** > **Connectors**.

Step 3 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit or delete a connector: Click **More** ( ), then click **Edit** or **Delete** at the end of the row.

Step 4 Enter the following information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
Pull Interval	(Default 30 seconds.) Interval at which IP mappings are retrieved from AWS.
Region	(Required.) Enter your AWS region code.
Access Key	(Required.) Enter your access key.
Secret Key	(Required.) Enter your secret key.

Step 5 Click **Save**.

Step 6 Make sure **Ok** is displayed in the Status column.

Amazon Web Services Security Groups Connector—About User Permissions and Imported Data

The Cisco Secure Dynamic Attributes Connector imports dynamic attributes from AWS to Secure Firewall Management Center for use in policies.

Minimum permissions required

The Cisco Secure Dynamic Attributes Connector requires a user at minimum with a policy that permits `ec2:DescribeTags`, `ec2:DescribeVpcs`, and `ec2:DescribeInstances` to be able to import dynamic attributes.

Create an AWS User with Minimal Permissions for the Cisco Secure Dynamic Attributes Connector

This task discusses how to set up a service account with minimum permissions to send dynamic attributes to Secure Firewall Management Center . For a list of these attributes, see Amazon Web Services Connector—About User Permissions and Imported Data, on page 1816.

Before you begin

You must already have set up your Amazon Web Services (AWS) account. For more information about doing that, see this article in the AWS documentation.

Procedure

Step 1 Log in to the AWS console as a user with the admin role.

Step 2 From the Dashboard, click **Security, Identity & Compliance** > **IAM**.

Step 3 Click **Access Management** > **Users**.

Step 4 Click **Add Users**.

Step 5 In the **User Name** field, enter a name to identify the user.

Step 6 Click **Access Key - Programmatic Access**.

Step 7 At the Set permissions page, click **Next** without granting the user access to anything. You can grant user access later.

Step 8 Add tags to the user if desired.

Step 9 Click **Create User**.

Step 10 Click **Download .csv** to download the user's key to your computer.

Note
This is the only opportunity you have to retrieve the user's key.

Step 11 Click **Close**.

Step 12 At the Identity and Access Management (IAM) page in the left column, click **Access Management** > **Policies**.

Step 13 Click **Create Policy**.

Step 14 On the Create Policy page, click **JSON**.

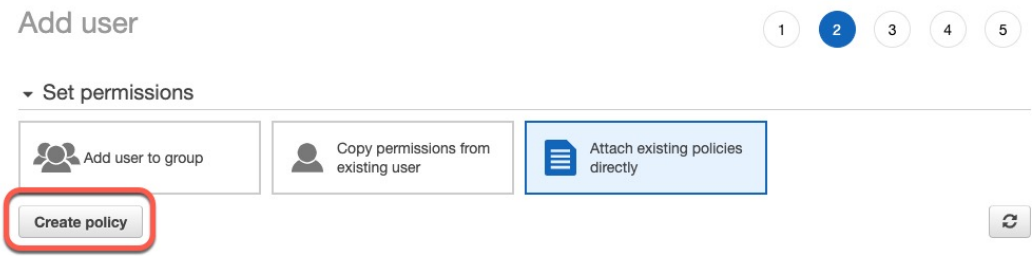

Step 15 Enter the following policy in the field:

```
{
 "Version": "2012-10-17",
 "Statement": [
  {
   "Effect": "Allow",
   "Action": [
    "ec2:DescribeTags",
    "ec2:DescribeInstances",
    "ec2:DescribeVpcs"
   ],
   "Resource": "*"
  }
```

```
  ]
}
```

Step 16 Click **Next**.

Step 17 Click **Review**.

Step 18 On the Review Policy page, enter the requested information and click **Create Policy**.

Step 19 On the Policies page, enter all or part of the policy name in the search field and press Enter.

Step 20 Click the policy you just created.

Step 21 Click **Actions** > **Attach**.

Step 22 If necessary, enter all or part of the user name in the search field and press Enter.

Step 23 Click **Attach Policy**.

What to do next

Create an AWS Connector, on page 1818.

Create an AWS Security Groups Connector

This task discusses how to configure a connector that sends AWS security groups data to the Secure Firewall Management Center for use in policies.

Before you begin

Do all of the following:

- Create AWS security groups as discussed in Work with security groups on the AWS documentation site.
- Create a user with at least the privileges discussed in Create an AWS User with Minimal Permissions for the Cisco Secure Dynamic Attributes Connector, on page 1817.

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Integration** > **Dynamic Attributes Connector** > **Connectors**.

Step 3 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit or delete a connector: Click **More** (⋮), then click **Edit** or **Delete** at the end of the row.

Step 4 Enter the following information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.

Value	Description
Pull Interval	(Default 30 seconds.) Interval at which IP mappings are retrieved from AWS. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
Region	(Required.) Enter your AWS region code.
AWS Access Key	(Required.) Enter your access key.
AWS Secret Key	(Required.) Enter your secret key.

Step 5 Click **Save**.

Step 6 Make sure **Ok** is displayed in the Status column.

Create an AWS Service Tags Connector

This topic discusses how to create a connector for Amazon Web Services (AWS) service tags to the Secure Firewall Management Center for use in policies.

For more information, see resources like the following on the AWS documentation site:

- What are tags?
- AWS IP address ranges
- Tagging your AWS resources
- Guidance for Tagging on AWS
- AWS service points

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Integration** > **Dynamic Attributes Connector** > **Connectors**.

Step 3 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit or delete a connector: Click **More** ( ), then click **Edit** or **Delete** at the end of the row.

Step 4 Enter the following information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
URL	(Required.) Do not change the URL unless advised to do so.

Step 5 Click **Save**.

Step 6 Make sure **Ok** is displayed in the Status column.

Azure Connector—About User Permissions and Imported Data

The Cisco Secure Dynamic Attributes Connector imports dynamic attributes from Azure to Secure Firewall Management Center for use in policies.

Dynamic attributes imported

We import the following dynamic attributes from Azure:

- *Tags*, key-value pairs associated with resources, resource groups, and subscriptions.

  For more information, see this page in the Microsoft documentation.

- *IP addresses* of virtual machines in Azure.

Minimum permissions required

The Cisco Secure Dynamic Attributes Connector requires a user at minimum with the **Reader** permission to be able to import dynamic attributes.

Create an Azure User with Minimal Permissions for the Cisco Secure Dynamic Attributes Connector

This task discusses how to set up a service account with minimum permissions to send dynamic attributes to Secure Firewall Management Center . For a list of these attributes, see Azure Connector—About User Permissions and Imported Data, on page 1823.

Before you begin

You must already have a Microsoft Azure account. To set one up, see this page on the Azure documentation site.

Procedure

Step 1 Log in to the Azure Portal as the owner of the subscription.

Step 2 Click **Azure Active Directory**.

Step 3 Find the instance of Azure Active Directory for the application you want to set up.

Step 4 Click **Add** > **App registration**.

Step 5 In the **Name** field, enter a name to identify this application.

Step 6 Enter other information on this page as required by your organization.

Step 7 Click **Register**.

Step 8 On the next page, write down or copy the Client ID (also referred to as *application ID*) and the tenant ID (also referred to as the *directory ID*).

A sample follows.

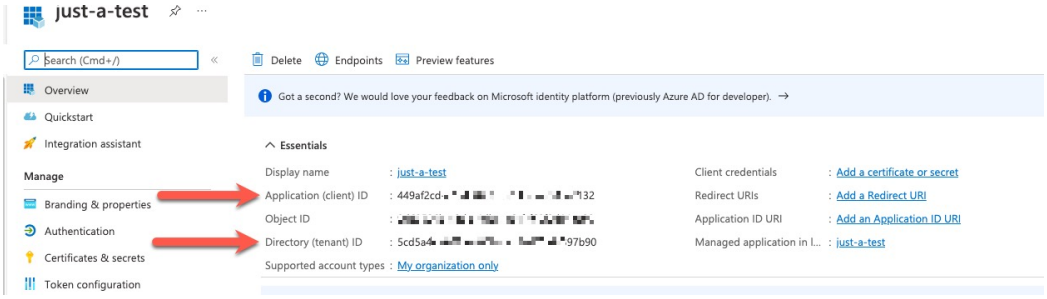

Step 9 Next to Client Credentials, click **Add a certificate or secret**.

Step 10 Click **New Client Secret**.

Step 11 Enter the requested information and click **Add**.

Step 12 Copy the value of the **Value** field to the clipboard. This value, *and not the* ***Secret ID***, is the client secret.

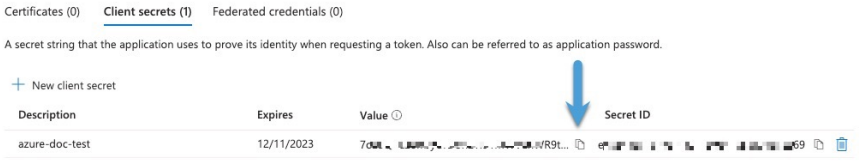

Step 13 Go back to the main Azure Portal page and click **Subscriptions**.

Step 14 Click the name of your subscription.

Step 15 Copy the subscription ID to the clipboard.

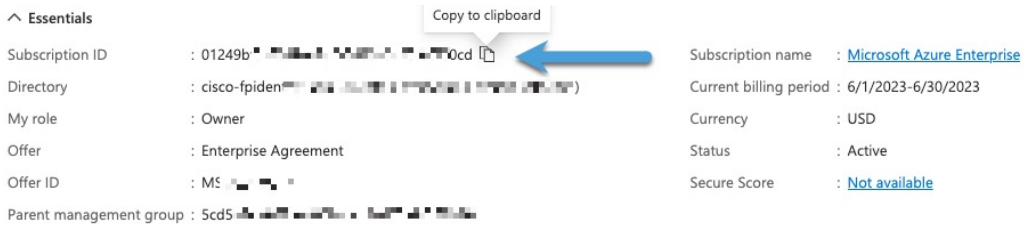

Step 16 Click **Access Control (IAM)**.

Step 17 Click **Add** > **Add role assignment**.

Step 18 Click **Reader** and click **Next**.

Step 19 Click **Select Members**.

Step 20 On the right side of the page, click the name of the app you registered and click **Select**.

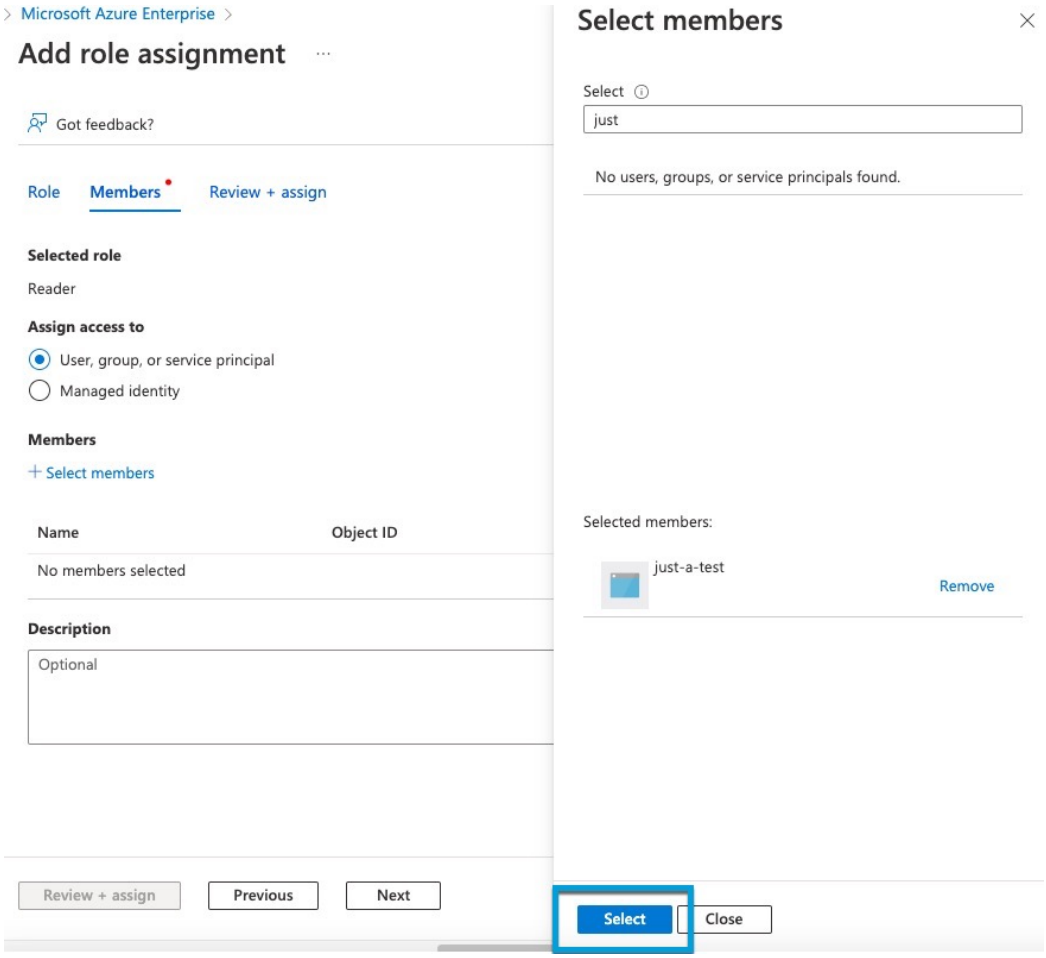

Step 21 Click **Review + Assign** and follow the prompts to complete the action.

What to do next

See Create an Azure Connector, on page 1825.

Create an Azure Connector

This task discusses how to create a connector to send data from Azure to Secure Firewall Management Center for use in policies.

Before you begin

Create an Azure user with at least the privileges discussed in Create an Azure User with Minimal Permissions for the Cisco Secure Dynamic Attributes Connector, on page 1823.

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Integration** > **Dynamic Attributes Connector** > **Connectors**.

Step 3 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit or delete a connector: Click **More** (⋮), then click **Edit** or **Delete** at the end of the row.

Step 4 Enter the following information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
Pull Interval	(Default 30 seconds.) Interval at which IP mappings are retrieved from Azure. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
Subscription Id	(Required.) Enter your Azure subscription ID.
Tenant Id	(Required.) Enter your tenant ID.
Client Id	(Required.) Enter your client ID.
Client Secret	(Required.) Enter your client secret.

Step 5 Click **Save**.

Step 6 Make sure **Ok** is displayed in the Status column.

Create an Azure Service Tags Connector

This topic discusses how to create a connector for Azure service tags to the Secure Firewall Management Center for use in policies. The IP addresses associated with these tags are updated every week by Microsoft.

For more information, see Virtual network service tags on Microsoft TechNet.

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Integration** > **Dynamic Attributes Connector** > **Connectors**.

Step 3 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit or delete a connector: Click **More** (), then click **Edit** or **Delete** at the end of the row.

Step 4 Enter the following information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
Pull Interval	(Default 30 seconds.) Interval at which IP mappings are retrieved from Azure. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
Subscription Id	(Required.) Enter your Azure subscription ID.
Tenant Id	(Required.) Enter your tenant ID.
Client Id	(Required.) Enter your client ID.
Client Secret	(Required.) Enter your client secret.

Step 5 Click **Save**.

Step 6 Make sure **Ok** is displayed in the Status column.

Create a Cisco Cyber Vision Connector

This task discusses how to send data from Cisco Cyber Vision to the Secure Firewall Management Center .

Before you begin

Cisco Cyber Vision must be reachable from the machine on which the dynamic attributes connector is running. You must know its IP address, port, and API key.

To find the API key in the Cyber Vision management console, click **Admin** > **API** > **Token**, then click **Show** to display the token and to copy the token to the clipboard.

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Integration** > **Dynamic Attributes Connector** > **Connectors**.

Step 3 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit or delete a connector: Click **More** ( ), then click **Edit** or **Delete** at the end of the row.

Step 4 Enter the following information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
Cyber Vision Prefix	Enter an alphanumeric string to identify dynamic objects from this Cyber Vision's IP address when objects are sent to Secure Firewall Management Center . If you have one Cyber Vision IP address, you can enter any value such as **1**.
Pull Interval	(Default 60 seconds.) Interval at which data mappings are retrieved from Cyber Vision. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
IP	(Required.) Enter the Cyber Vision IP address.
Port	(Required.) Enter the Cyber Vision listen port.
Token	(Required.) Enter the API token.

Step 5 Click **Test** and make sure the test succeeds before you save the connector.

Step 6 Click **Save**.

Step 7 Make sure **Ok** is displayed in the Status column.

Create a Generic Text Connector

This task discusses how to create an ad hoc list of IP addresses you maintain manually and retrieve at an interval you select (30 seconds by default). You can update the list of addresses anytime you want.

Before you begin

Create text files with IP addresses and put it on a web server that is accessible from the Secure Firewall Management Center . IP addresses can include CIDR notation. The text file must have only one IP address per line.

For example, you might have a list of IP addresses for an "allow list" in access control rules and another list of IP addresses for a "block list" in access control rules.

You can specify up to 10,000 IP addresses per text file.

Note *Do not* include a scheme (**`http://`** or **`https://`**) in your IP addresses.

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Integration** > **Dynamic Attributes Connector** > **Connectors**.

Step 3 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit or delete a connector: Click **More** ( ), then click **Edit** or **Delete** at the end of the row.

Step 4 Enter the following information:

Item	Description
Name	Enter a name to identify the connector.
Description	(Optional.) Enter a description
Pull Interval	Change the frequency, in seconds, at which the dynamic attributes connector retrieves IP addresses from the text file. The default is 30 seconds. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
URLs	Enter each URL from which to retrieve IP addresses, one URL per line.
Add another URL	(Optional.) Click the link to add more URLs to an existing list.
Certificate	(Optional.) If a certificate chain is required for a secure connection to the web server, you have the following options: • Click **Get Certificate** > **Fetch** to automatically fetch the certificate or, if that is not possible, get the certificate manually as discussed in Manually Get a Certificate Authority (CA) Chain, on page 1842. • Click **Get Certificate** > **Browse from file** to upload a certificate chain you downloaded previously.

Step 5 Click **Test** and make sure the test succeeds before you save the connector.

Step 6 Click **Save**.

Step 7 Make sure **Ok** is displayed in the Status column.

Create a GitHub Connector

This section discusses how to create a GitHub connector that sends data to the Secure Firewall Management Center for use in policies. The IP addresses associated with these tags are maintained by GitHub. You do not have to create a dynamic attributes filters.

For more information, see About GitHub's IP addresses.

Note Do not change the URL because doing so will fail to retrieve any IP addresses.

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Integration** > **Dynamic Attributes Connector** > **Connectors**.

Step 3 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit or delete a connector: Click **More** ( ), then click **Edit** or **Delete** at the end of the row.

Step 4 Enter a **Name** and an optional description.

Step 5 (Optional.) In the **Pull Interval** field, change the frequency, in seconds, at which the dynamic attributes connector retrieves IP addresses from GitHub. The default is 21,600 seconds (6 hours).

Step 6 Click **Save**.

Step 7 Make sure **Ok** is displayed in the Status column.

Google Cloud Connector—About User Permissions and Imported Data

The Cisco Secure Dynamic Attributes Connector imports dynamic attributes from Google Cloud to Secure Firewall Management Center for use in policies.

Dynamic attributes imported

We import the following dynamic attributes from Google Cloud:

- *Labels*, key-value pairs you can use to organize your Google Cloud resources.

  For more information, see Creating and Managing Labels in the Google Cloud documentation.

- *Network tags*, key-value pairs associated with an organization, folder, or project.

  For more information, see Creating and Managing Tags in the Google Cloud documentation.

- *IP addresses* of virtual machines in Google Cloud.

Minimum permissions required

The Cisco Secure Dynamic Attributes Connector requires a user at minimum with the **Basic** > **Viewer** permission to be able to import dynamic attributes.

Create a Google Cloud User with Minimal Permissions for the Cisco Secure Dynamic Attributes Connector

This task discusses how to set up a service account with minimum permissions to send dynamic attributes to Secure Firewall Management Center . For a list of these attributes, see Google Cloud Connector—About User Permissions and Imported Data, on page 1830.

Before you begin

You must already have set up your Google Cloud account. For more information about doing that, see Setting Up Your Environment in the Google Cloud documentation.

Procedure

Step 1 Log in to your Google Cloud account as a user with the owner role.

Step 2 Click **IAM & Admin** > **Service Accounts** > **Create Service Account**.

Step 3 Enter the following information:

- **Service account name**: A name to identify this account; for example, **CSDAC**.
- **Service account ID**: Should be populated with a unique value after you enter the service account name.
- **Service account description**: Enter an optional description.

For more information about service accounts, see Understanding Service Accounts in the Google Cloud documentation.

Step 4 Click **Create and Continue**.

Step 5 Follow the prompts on your screen until the Grant users access to this service account section is displayed.

Step 6 Grant the user the **Basic** > **Viewer** role.

Step 7 Click **Done**.

A list of service accounts is displayed.

Step 8 Click **More** (⋮) at the end of the row of the service account you created.

Step 9 Click **Manage Keys**.

Step 10 Click **Add Key** > **Create New Key**.

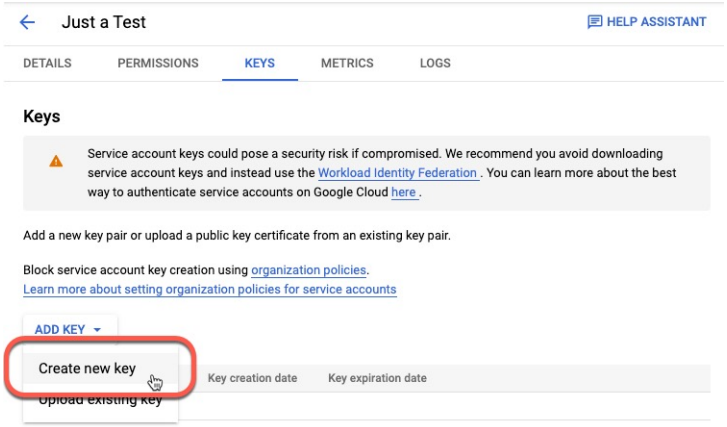

Step 11 Click **JSON**.

Step 12 Click **Create**.

The JSON key is downloaded to your computer.

Step 13 Keep the key handy when you configure the GCP connector.

What to do next

See Create a Google Cloud Connector, on page 1832.

Create a Google Cloud Connector

Before you begin

Have your Google Cloud JSON-formatted service account data ready; it's required to set up the connector.

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Integration** > **Dynamic Attributes Connector** > **Connectors**.

Step 3 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit or delete a connector: Click **More** ( ), then click **Edit** or **Delete** at the end of the row.

Step 4 Enter the following information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.

Value	Description
Pull Interval	(Default 30 seconds.) Interval at which IP mappings are retrieved from AWS. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
GCP region	(Required.) Enter the GCP region in which your Google Cloud is located. For more information, see Regions and Zones in the Google Cloud documentation.
Service account	Paste the JSON code for your Google Cloud service account.

Step 5 Click **Save**.

Step 6 Make sure **Ok** is displayed in the Status column.

Create an Office 365 Connector

This task discusses how to create a connector for Office 365 tags to send data to the Secure Firewall Management Center for use in policies. The IP addresses associated with these tags are updated every week by Microsoft. You do not have to create a dynamic attributes filter to use the data.

For more information, see Office 365 URLs and IP address ranges on docs.microsoft.com.

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Integration** > **Dynamic Attributes Connector** > **Connectors**.

Step 3 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit or delete a connector: Click **More** ( ), then click **Edit** or **Delete** at the end of the row.

Step 4 Enter the following information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
Pull Interval	(Default 30 seconds.) Interval at which IP mappings are retrieved from Azure. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.

Value	Description
Base API URL	(Required.) Enter the URL from which to retrieve Office 365 information, if it's different from the default. For more information, see Office 365 IP Address and URL web service on the Microsoft documentation site.
Instance name	(Required.) From the list, click an instance name. For more information, see Office 365 IP Address and URL web service on the Microsoft documentation site.
Disable optional IPs	(Required.) Enter **true** or **false**.

Step 5 Click **Save**.

Step 6 Make sure **Ok** is displayed in the Status column.

vCenter Connector—About User Permissions and Imported Data

The Cisco Secure Dynamic Attributes Connector imports dynamic attributes from vCenter to Secure Firewall Management Center for use in policies.

Dynamic attributes imported

We import the following dynamic attributes from vCenter:

- *Operating system*
- *MAC address*
- *IP addresses*
- *NSX tags*

Minimum permissions required

The Cisco Secure Dynamic Attributes Connector requires a user at minimum with the **Read Only** permission to be able to import dynamic attributes.

Create a vCenter User with Minimal Permissions for the Cisco Secure Dynamic Attributes Connector

This task discusses how to set up a service account with minimum permissions to send dynamic attributes to Secure Firewall Management Center . For a list of these attributes, see vCenter Connector—About User Permissions and Imported Data, on page 1834.

Before you begin

You must already have set up your vCenter Server account. For more information about doing that, see About vCenter Server Installation and Setup in the vCenter documentation.

Procedure

Step 1 Log into vCenter as an administrator.

Step 2 Click **Menu** > **Administration**.

Step 3 In the left pane, click **Single Sign On** > **Users and Groups**.

Step 4 From the **Domain** list, click the name of a domain to add the user.

Step 5 Click **Add User**.

Step 6 Enter the requested information and click **Add**.

Step 7 In the left pane, click **Access Control** > **Global Permissions**.

Step 8 Click **Add**(+).

Step 9 From the **User** field, click the name of the vCenter domain in which you created the user.

Step 10 In the search field, enter part of the user's name.

Step 11 From the **Role** list, click **Read-only**.

Step 12 Select the **Propagate to children** check box.

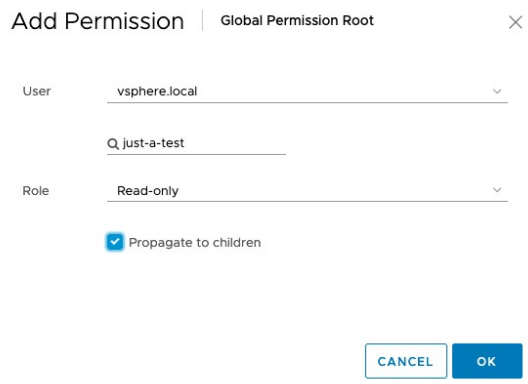

Step 13 Click **OK**.

What to do next

See Create a vCenter Connector, on page 1835.

Create a vCenter Connector

This task discusses how to create a connector for VMware vCenter to send data to the Secure Firewall Management Center for use in policies.

Before you begin

If you use non-trusted certificates to communicate with vCenter, see Manually Get a Certificate Authority (CA) Chain, on page 1842.

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Integration** > **Dynamic Attributes Connector** > **Connectors**.

Step 3 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit or delete a connector: Click **More** (⋮), then click **Edit** or **Delete** at the end of the row.

Step 4 Enter the following information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Enter an optional description.
Pull Interval	(Default 30 seconds.) Interval at which IP mappings are retrieved from vCenter. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
Host	(Required.) Enter any of the following: • vCenter's fully qualified host name • vCenter's IP address • (Optional.) A port *Do not* enter a scheme (such as **`https://`**) or trailing slash. For example, **`myvcenter.example.com`** or **`192.0.2.100:9090`**
User	(Required.) Enter the user name of a user with the Read-only role at minimum. User names are case-sensitive.
Password	(Required.) Enter the user's password.
NSX IP	If you use vCenter Network Security Visualization (NSX), enter its IP address.
NSX User	Enter the user name of an NSX user with the Auditor role at minimum.
NSX Type	Enter **NSX-T**.
NSX Password	Enter the NSX user's password.

Value	Description
vCenter Certificate	You have the following options: • Paste the certificate authority (CA) chain you got as discussed in Manually Get a Certificate Authority (CA) Chain, on page 1842. • Click **Fetch** to automatically fetch the certificate or, if that is not possible, get the certificate manually as discussed in Manually Get a Certificate Authority (CA) Chain, on page 1842. • Click **Get Certificate** > **Fetch** to automatically fetch the certificate or, if that is not possible, get the certificate manually as discussed in Manually Get a Certificate Authority (CA) Chain, on page 1842. • Click **Get Certificate** > **Browse from file** to upload a certificate chain you downloaded previously.

Following is an example of successfully fetching a certificate chain:

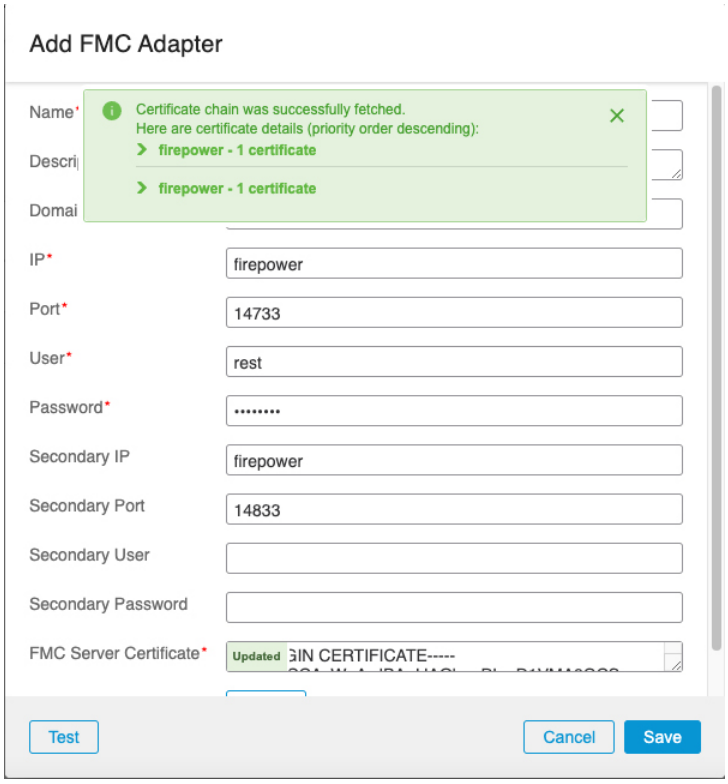

Expanding the certificate CA chain at the top of the dialog box displays the certificates similar to the following.

If it's not possible to fetch the certificate this way, you can get the certificate chain manually as discussed in Manually Get a Certificate Authority (CA) Chain, on page 1842.

Step 5 Click **Save**.

Create a Webex Connector

This section discusses how to create a Webex connector that sends data to the Secure Firewall Management Center for use in policies. The IP addresses associated with these tags are maintained by Webex. You do not have to create a dynamic attributes filters.

For more information, see Port Reference for Webex Calling.

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Integration** > **Dynamic Attributes Connector** > **Connectors**.

Step 3 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit or delete a connector: Click **More** ( ), then click **Edit** or **Delete** at the end of the row.

Step 4 Enter the following information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
Pull Interval	(Default 30 seconds.) Interval at which IP mappings are retrieved from Webex. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
Provider Reserved IPs	(Required.) (Required.) Slide to enabled to retrieve any reserved IP addresses.

Step 5 Click **Test** and make sure the test succeeds before you save the connector.

Step 6 Click **Save**.

Step 7 Make sure **Ok** is displayed in the Status column.

Create a Zoom Connector

This section discusses how to create a Zoom connector that sends data to the Secure Firewall Management Center for use in policies. The IP addresses associated with these tags are maintained by Zoom. You do not have to create a dynamic attributes filters.

For more information, see Zoom network firewall or proxy server settings.

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Integration** > **Dynamic Attributes Connector** > **Connectors**.

Step 3 Do any of the following:

- Add a new connector: click Add icon ( ), then click the name of the connector.
- Edit or delete a connector: Click **More** (⋮), then click **Edit** or **Delete** at the end of the row.

Step 4 Enter the following information.

Value	Description
Name	(Required.) Enter a name to uniquely identify this connector.
Description	Optional description.
Pull Interval	(Default 30 seconds.) Interval at which IP mappings are retrieved from Zoom. The minimum value for **Pull Interval** is 1 second. You can set the maximum to any value you want. We recommend against setting the minimum to a low value because it can generate a lot of traffic, and, when applicable, can result in your being billed for the traffic.
Provider Reserved IPs	(Required.) Slide to enabled to retrieve any reserved IP addresses.

Step 5 Click **Test** and make sure the test succeeds before you save the connector.

Step 6 Click **Save**.

Step 7 Make sure **Ok** is displayed in the Status column.

Create Dynamic Attributes Filters

Dynamic attributes filters that you define using the Cisco Secure Dynamic Attributes Connector are exposed in the Secure Firewall Management Center as dynamic objects that can be used in access control policies. For example, restrict access to an AWS server for the Finance Department to only members of the Finance group defined in Microsoft Active Directory.

Note You cannot create dynamic attributes filters for Generic Text, Office 365, Azure Service Tags, Webex, or Zoom. These types of cloud objects provide their own IP addresses.

For more information about access control rules, see Create Access Control Rules Using Dynamic Attributes Filters, on page 1846.

Before you begin

Create a Connector, on page 1816

Procedure

Step 1 Log in to Secure Firewall Management Center.

Step 2 Click **Integration** > **Dynamic Attributes Connector**.

Step 3 Click **Dynamic Attributes Filters**.

- Add a new filter: click **Add** ( ).
- Edit or delete a filter: Click **More** ( ), then click **Edit** or **Delete** at the end of the row.

Step 4 Enter the following information.

Item	Description
Name	Unique name to identify the dynamic filter (as a dynamic object) in a policy and in the Secure Firewall Management Center Object Manager (**External Attributes** > **Dynamic Object**).
Connector	From the list, click the name of a connector to use.
Query	Click Add .

Step 5 To add or edit a query, enter the following information.

Item	Description
Key	Click a key from the list. Keys are fetched from the connector.
Operation	Click one of the following: • **Equals** to exactly match the key to the value. • **Contains** to match the key to the value if any part of the value matches.

Item	Description
Values	Click either **Any** or **All** and click one or more values from the list. Click **Add another value** to add values to your query.

Step 6 Click **Show Preview** to display a list of networks or IP addresses returned by your query.

Step 7 When you're finished, click **Save**.

Step 8 (Optional.) Verify the dynamic object in the Secure Firewall Management Center .

a) Log in to the Secure Firewall Management Center as a user with the Network Admin role at minimum.

b) Click **Objects** > **Object Management**.

c) In the left pane, click **External Attributes** > **Dynamic Object**.

The dynamic attribute query you created should be displayed as a dynamic object.

Dynamic Attribute Filter Examples

This topic provides some examples of setting up dynamic attribute filters.

Examples: vCenter

The following example shows one criterion: a VLAN.

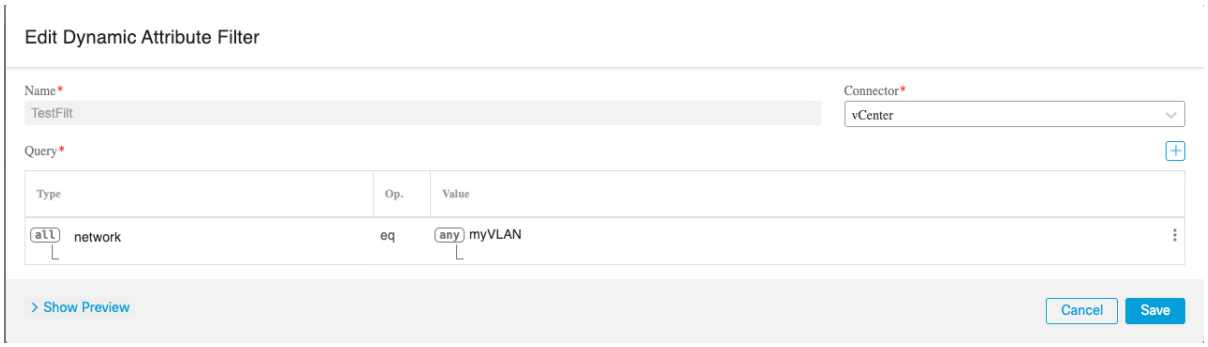

The following example shows three criteria that are joined with OR: the query matches any of three hosts.

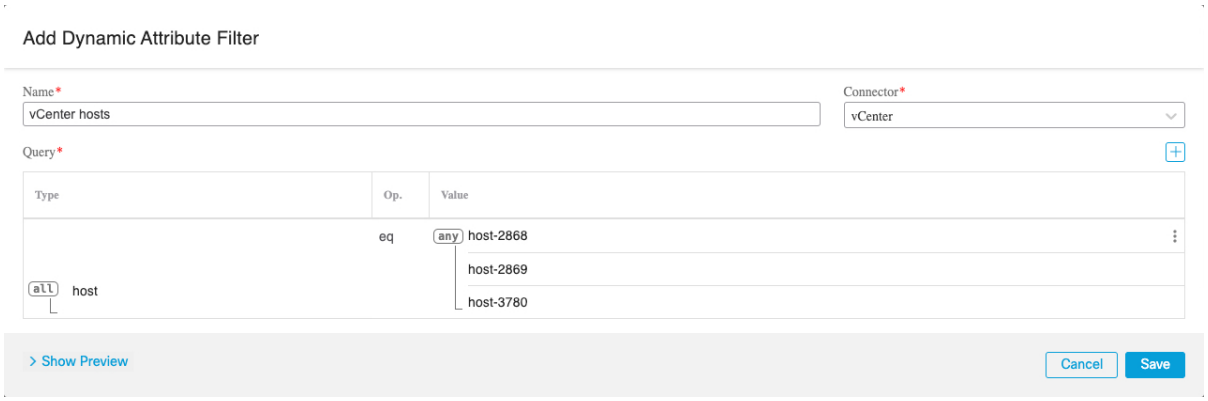

Example: Azure

The following example shows one criterion: a server tagged as a Finance app.

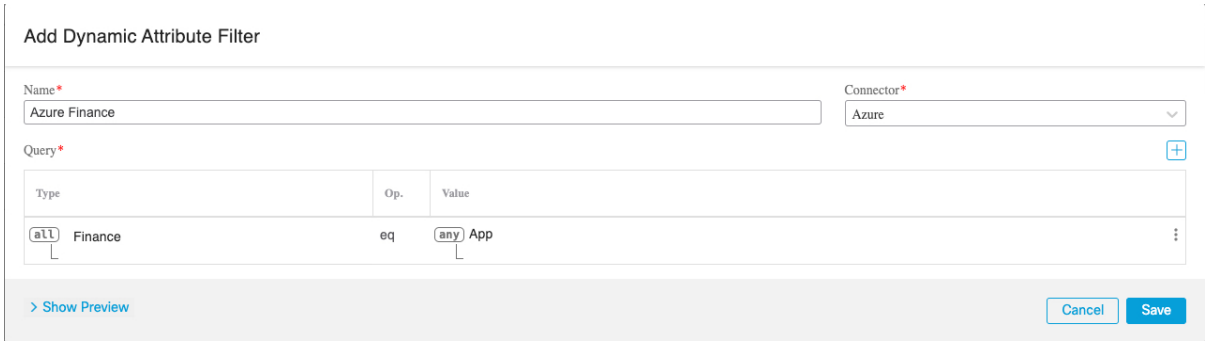

Example: AWS

The following example shows one criterion: a FinanceApp with a value of 1.

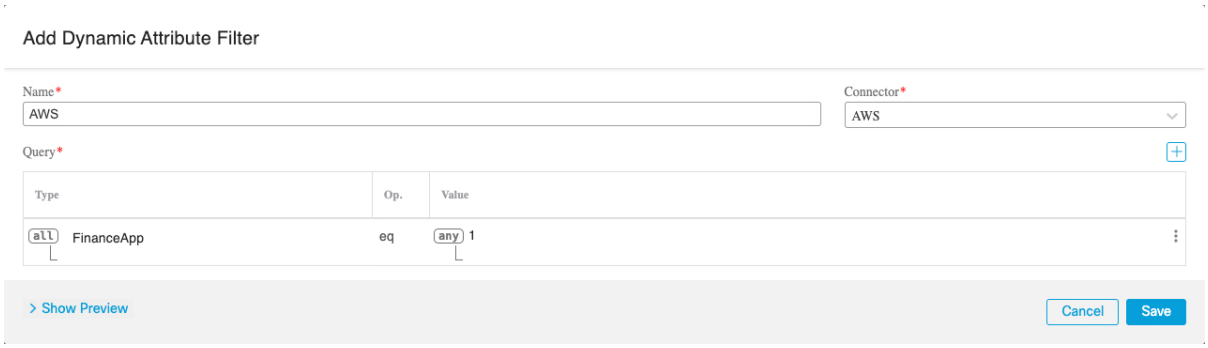

Manually Get a Certificate Authority (CA) Chain

In the event you cannot automatically fetch the certificate authority chain, use one of the following browser-specific procedures to get a certificate chain used to connect securely to vCenter or Firewall Management Center.

The *certificate chain* is the root certificate and all subordinate certificates.

You can optionally use one of these procedures to connect to the following:

- vCenter or NSX
- Firewall Management Center

Get a Certificate Chain—Mac (Chrome and Firefox)

Use this procedure to get a certificate chain using the Chrome and Firefox browsers on Mac OS.

1. Open a Terminal window.
2. Enter the following command.

```
security verify-cert -P url[:port]
```

where *url* is the URL (including scheme) to vCenter or Firewall Management Center. For example:

```
security verify-cert -P https://myvcenter.example.com
```

If you access vCenter or Firewall Management Center using NAT or PAT, you can add a port as follows:

```
security verify-cert -P https://myvcenter.example.com:12345
```

3. Save the entire certificate chain to a plaintext file.
   - *Include* all `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` delimiters.
   - *Exclude* any extraneous text (for example, the name of the certificate and any text contained in angle brackets (< and >) as well as the angle brackets themselves.
4. Repeat these tasks for vCenter Firewall Management Center.

Get a Certificate Chain—Windows Chrome

Use this procedure to get a certificate chain using the Chrome browser on Windows.

1. Log in to vCenter or Firewall Management Center using Chrome.
2. In the browser address bar, click the lock to the left of the host name.
3. Click **Certificate**.
4. Click the **Certification Path** tab.
5. Click the top (that is, first) certificate in the chain.
6. Click **View Certificate**.
7. Click the **Details** tab.
8. Click **Copy to File**.
9. Follow the prompts to create a CER-formatted certificate file that includes the entire certificate chain.

   When you're prompted to choose an export file format, click **Base 64-Encoded X.509 (.CER)** as the following figure shows.

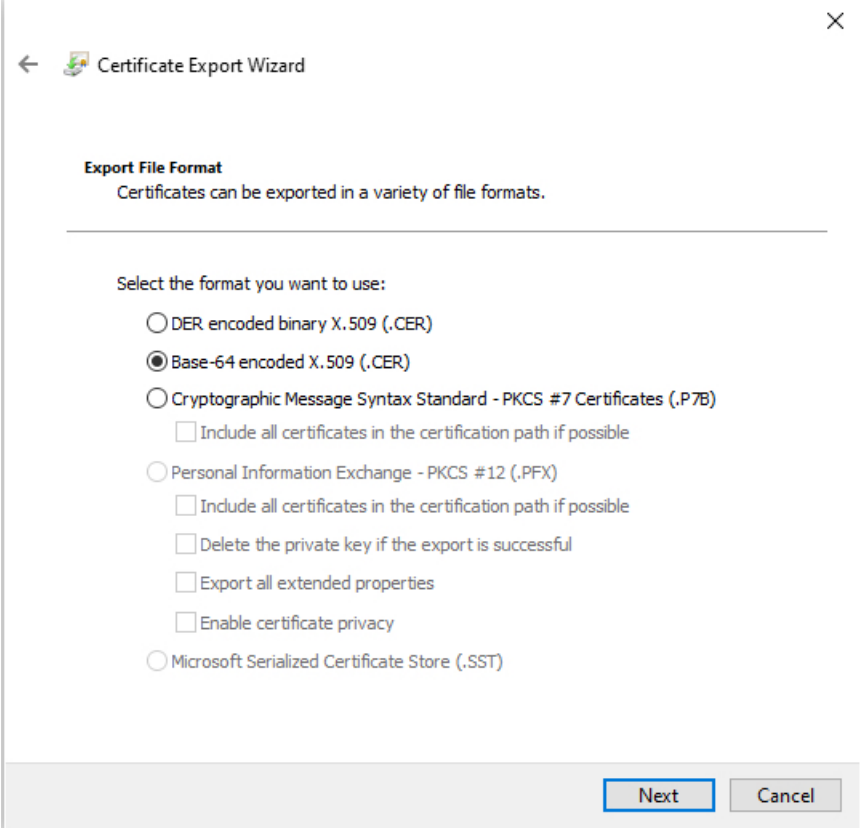

10. Follow the prompts to complete the export.
11. Open the certificate in a text editor.
12. Repeat the process for all certificates in the chain.

    You must paste each certificate in the text editor in order, first to last.
13. Repeat these tasks for vCenter or Firewall Management Center.

Get a Certificate Chain—Windows Firefox

Use the following procedure to get a certificate chain for the Firefox browser on either Windows or Mac OS.

1. Log in to vCenter or Firewall Management Center. using Firefox.
2. Click the lock to the left of the host name.
3. Click the right arrow (**Show connection details**). The following figure shows an example.

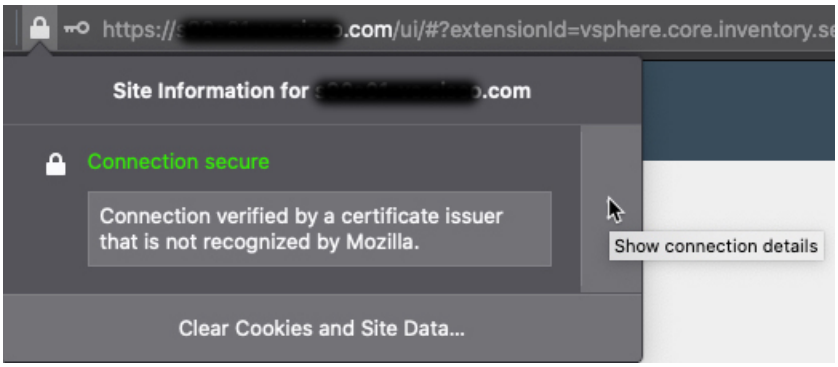

4. Click **More Information**.
5. Click **View Certificate**.
6. If the resulting dialog box has tab pages, click the tab page corresponding to the top-level CA.
7. Scroll to the Miscellaneous section.
8. Click **PEM (chain)** in the Download row. The following figure shows an example.

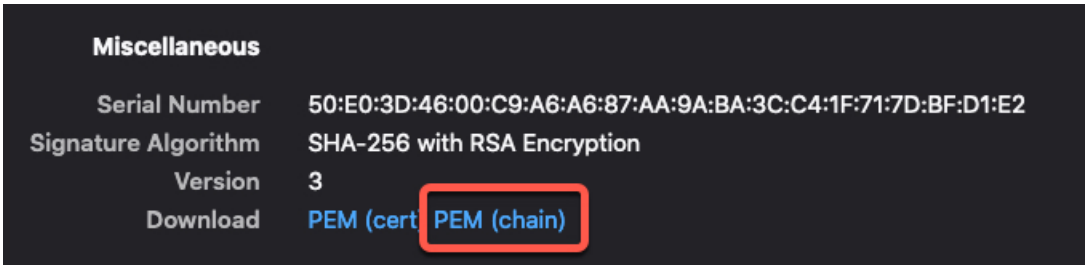

9. Save the file.
10. Repeat these tasks for vCenter or Firewall Management Center.

Use Dynamic Objects in Access Control Policies

The dynamic attributes connector enables you to configure dynamic filters, seen in the Secure Firewall Management Center as dynamic objects, in access control rules.

About Dynamic Objects in Access Control Rules

A *dynamic object* is automatically pushed from the dynamic attributes connector to the Secure Firewall Manager after you create connectors and save a dynamic attributes filter on the connector.

You can use these dynamic objects on the access control rule's **Dynamic Attributes** tab page, similarly to the way you used Security Group Tags (SGTs). You can add dynamic objects as source or destination attributes; for example, in an access control block rule, you can add a Finance dynamic object as a destination attribute to block access to Finance servers by whatever objects match the other criteria in the rule.

Note You cannot create dynamic attributes filters for Generic Text, Office 365, Azure Service Tags, Webex, or Zoom. These types of cloud objects provide their own IP addresses.

Dynamic Attributes Rule Conditions

Dynamic attributes include the following:

- (Source or destination.) Dynamic objects (such as from the Cisco Secure Dynamic Attributes Connector)

  The dynamic attributes connector enables you to collect data (such as networks and IP addresses) from cloud providers and send it to the Secure Firewall Management Center so it can be used in access control rules.

  For more information about the dynamic attributes connector, see About the Cisco Secure Dynamic Attributes Connector, on page 1803.

- (Source only.) Location IP objects, defined by Cisco ISE
- (Source only.) Device type objects, defined by Cisco ISE (also referred to as endpoint profile objects)

Dynamic attributes can be used as source criteria and destination criteria in access control rules. Use the following guidelines:

- Objects of different types are ANDd together
- Objects of a similar type are ORd together

For example, if you choose source destination criteria SGT 1, SGT 2, and device type 1; the rule is matched if device type 1 is detected on either SGT 1 or SGT 2. As another example, if you select both a security group tag, and a dynamic object that lists IP addresses, the rule matches if traffic with the tag originates from (or is destined to) one of those IP addresses.

Create Access Control Rules Using Dynamic Attributes Filters

This topic discusses how to create access control rules using dynamic objects (these dynamic objects are named after the dynamic attributes filters you created previously).

Before you begin

Create dynamic attributes filters as discussed in Create Dynamic Attributes Filters, on page 1839.

Note You cannot create dynamic attributes filters for Generic Text, Office 365, Azure Service Tags, Webex, or Zoom. These types of cloud objects provide their own IP addresses.

Procedure

Step 1 Log in to the Secure Firewall Management Center

Step 2 Click **Policies** > **Access Control heading** > **Access Control**.

Step 3 Click **Edit** ( ) next to an access control policy.

Step 4 Click **Add Rule**.

Step 5 Click the **Dynamic Attributes** tab.

Step 6 In the Available Attributes section, from the list, click **Dynamic Objects**.

The following figure shows an example.

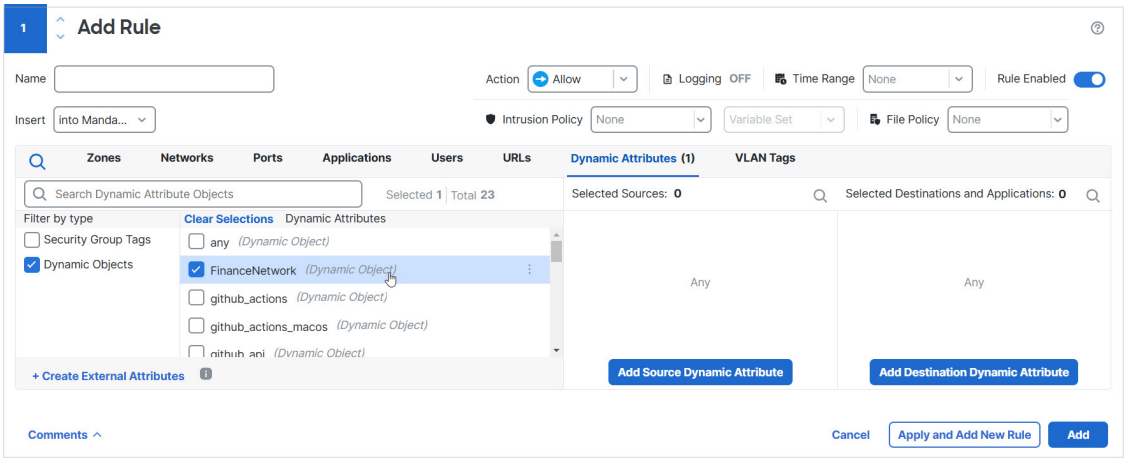

The preceding example shows a dynamic object named `APIC Dynamic Attribute` that corresponds to the dynamic attribute filter created in the Cisco Secure Dynamic Attributes Connector.

Step 7 Add the desired object to source or destination attributes.

Step 8 Add other conditions to the rule if desired.

What to do next

See Dynamic Attributes Rule Conditions, on page 1846.

Disable the Cisco Secure Dynamic Attributes Connector

If you no longer wish to collect dynamic objects from cloud sources, you can disable the Cisco Secure Dynamic Attributes Connector in the Secure Firewall Management Center as discussed in the following task.

Procedure

Step 1 Log in to the Secure Firewall Management Center if you have not done so already.

Step 2 Click **Integration** > **Dynamic Attributes Connector**.

Step 3 Slide to **Disabled**.

Troubleshoot Using the Command Line

To assist you with advanced troubleshooting and working with Cisco TAC, we provide the following troubleshooting tools. To use these tools, log in as any user to the Ubuntu host on which the dynamic attributes connector is running.

Check container status

To check the status of the dynamic attributes connector Docker containers, enter the following commands:

```
cd /usr/local/sf/csdac
sudo ./muster-cli status
```

Sample output follows:

```
================================================ CORE SERVICES
================================================
          Name                          Command               State                    Ports

-------------------------------------------------------------------------------------------------------------
muster-bee                 /bin/sh -c /app/bee              Up
127.0.0.1:15050->50050/tcp, 50443/tcp
muster-envoy              /docker-entrypoint.sh runs ...   Up      127.0.0.1:6443->8443/tcp

muster-local-fmc-adapter   ./docker-entrypoint.sh run ...   Up

muster-ui-backend          ./docker-entrypoint.sh run ...   Up      50031/tcp

muster-user-analysis       ./docker-entrypoint.sh run ...   Up      50070/tcp

============================ CONNECTORS AND ADAPTERS ============================
             Name                                Command               State     Ports
-------------------------------------------------------------------------------------
muster-connector-o365.1.muster   ./docker-entrypoint.sh run ...   Up      50070/tcp
```

Stop, start, or restart the Cisco Secure Dynamic Attributes Connector Docker containers

If the **./muster-cli status** indicates containers are down or to restart containers in the event of issues, you can enter the following commands:

Stop and restart:

```
cd ~/csdac/app
sudo ./muster-cli stop
sudo ./muster-cli start
```

Start only:

```
cd ~/csdac/app
sudo ./muster-cli start
```

Enable application debug logging and generate troubleshoot files

If advised to do so by Cisco TAC, enable debug logging and generate troubleshoot files as follows:

```
cd ~/csdac/app
sudo ./muster-cli debug-on
sudo ./muster-cli ts-gen
```

The troubleshoot file name is **ts-bundle-***timestamp***.tar** and is created in the same directory.

The following table shows the location of troubleshoot files and logs in the troubleshoot file.

Location	What it contains
/csdac/app/ts-bundle-*timestamp***/info**	`etcd` database contents
/csdac/app/ts-bundle-*timestamp***/logs**	Container log files
/csdac/app/ts-bundle-*timestamp***/status.log**	Container status, versions, and image status

Enable debugging for a container

You can optionally enable debugging for individual containers if you first get the name of the container as follows:

```
cd /usr/local/sf/csdac
sudo ./muster-cli versions
```

Sample output follows:

```
CSDAC version: 1.0.0
CONTAINERS VERSIONS
CONTAINER                       | APP VERSION          | COMMIT
================================================================================
muster-bee                      | fmc7.4-13            |
944d50c6c384567693d6ecc5a31420de57f6ce2f
muster-envoy                    | fmc7.4-25            |
5e5f6d83164a4acbef5b106aa39e2e3f68fa738f
muster-local-fmc-adapter        | fmc7.4-17            |
c5902f818baa8e27d7c0b8027490dcacc28c0168
muster-ui-backend               | fmc7.4-64            |
165a1f5f0d763aa75829a30b5ffbddf0012682b6
muster-user-analysis            | fmc7.4-43            |
63cd64e29a92599908c3eb684d91e9f685d8c740
muster-connector-o365.1.muster | fmc7.4-8             |
28f075d315c8867f667b828970c9fbad35fa89cc
```

To enable debugging for the Office 365 connector, for example, enter the following command.

```
sudo ./muster-cli container-debug-on muster-connector-o365.1.muster
```

To disable debugging for that connector, enter the following command.

```
sudo ./muster-cli container-debug-off muster-connector-o365.1.muster
```

Verify dynamic objects on the Secure Firewall Management Center

To verify your connectors are creating objects on the Secure Firewall Management Center , you can use the following command on the Secure Firewall Management Center as an administrator:

```
sudo tail f /var/opt/CSCOpx/MDC/log/operation/usmsharedsvcs.log
```

Example: Successful object creation

```
26-Aug-2021 12:41:35.912,[INFO],(DefenseCenterServiceImpl.java:1442)
com.cisco.nm.vms.api.dc.DefenseCenterServiceImpl, ajp-nio-127.0.0.1-9009-exec-10
** REST Request [ CSM ]
** ID : 18b25356-fd6b-4cc4-8d27-bbccb52a6275
```

```
** URL: POST /audit
{
  "version": "7.1.0",
  "requestId": "18b25356-fd6b-4cc4-8d27-bbccb52a6275",
  "data": {
    "userName": "csdac-centos7",
    "subsystem": "API",
    "message": "POST
https://myfmc.example.com/api/fmc_config/v1/domain/e276abec-e0f2-11e3-8169-6d9ed49b625f
/object/dynamicobjects Created (201) - The request has been fulfilled and resulted in a new
 resource being created",
    "sourceIP": "192.0.2.103",
    "domainUuid": "e276abec-e0f2-11e3-8169-6d9ed49b625f",
    "time": "1629981695431"
  },
  "deleteList": []
}
```

Troubleshoot Using the Firewall Management Center

This task discusses how to generate troubleshoot files for the Secure Firewall Management Center.

Before you begin

For complete details about troubleshooting, see the troubleshooting chapter in the Cisco Secure Firewall Management Center Administration Guide.

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **System** (⚙) > **Health** > **Monitor**.

Step 3 In the left pane, click **Firewall Management Center**.

Step 4 At the top, click **System & Troubleshooting Details**.

Step 5 Click **Generate Troubleshooting Files**.

Step 6 Provide the files to Cisco TAC or to your Beta coordinator.

Manually Get a Certificate Authority (CA) Chain

In the event you cannot automatically fetch the certificate authority chain, use one of the following browser-specific procedures to get a certificate chain used to connect securely to vCenter or Firewall Management Center.

The *certificate chain* is the root certificate and all subordinate certificates.

You can optionally use one of these procedures to connect to the following:

- vCenter or NSX
- Firewall Management Center

Get a Certificate Chain—Mac (Chrome and Firefox)

Use this procedure to get a certificate chain using the Chrome and Firefox browsers on Mac OS.

1. Open a Terminal window.
2. Enter the following command.

   ```
   security verify-cert -P url[:port]
   ```

   where *url* is the URL (including scheme) to vCenter or Firewall Management Center. For example:

   ```
   security verify-cert -P https://myvcenter.example.com
   ```

   If you access vCenter or Firewall Management Center using NAT or PAT, you can add a port as follows:

   ```
   security verify-cert -P https://myvcenter.example.com:12345
   ```

3. Save the entire certificate chain to a plaintext file.
   - *Include* all `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` delimiters.
   - *Exclude* any extraneous text (for example, the name of the certificate and any text contained in angle brackets (< and >) as well as the angle brackets themselves.
4. Repeat these tasks for vCenter Firewall Management Center.

Get a Certificate Chain—Windows Chrome

Use this procedure to get a certificate chain using the Chrome browser on Windows.

1. Log in to vCenter or Firewall Management Center using Chrome.
2. In the browser address bar, click the lock to the left of the host name.
3. Click **Certificate**.
4. Click the **Certification Path** tab.
5. Click the top (that is, first) certificate in the chain.
6. Click **View Certificate**.
7. Click the **Details** tab.
8. Click **Copy to File**.
9. Follow the prompts to create a CER-formatted certificate file that includes the entire certificate chain.

   When you're prompted to choose an export file format, click **Base 64-Encoded X.509 (.CER)** as the following figure shows.

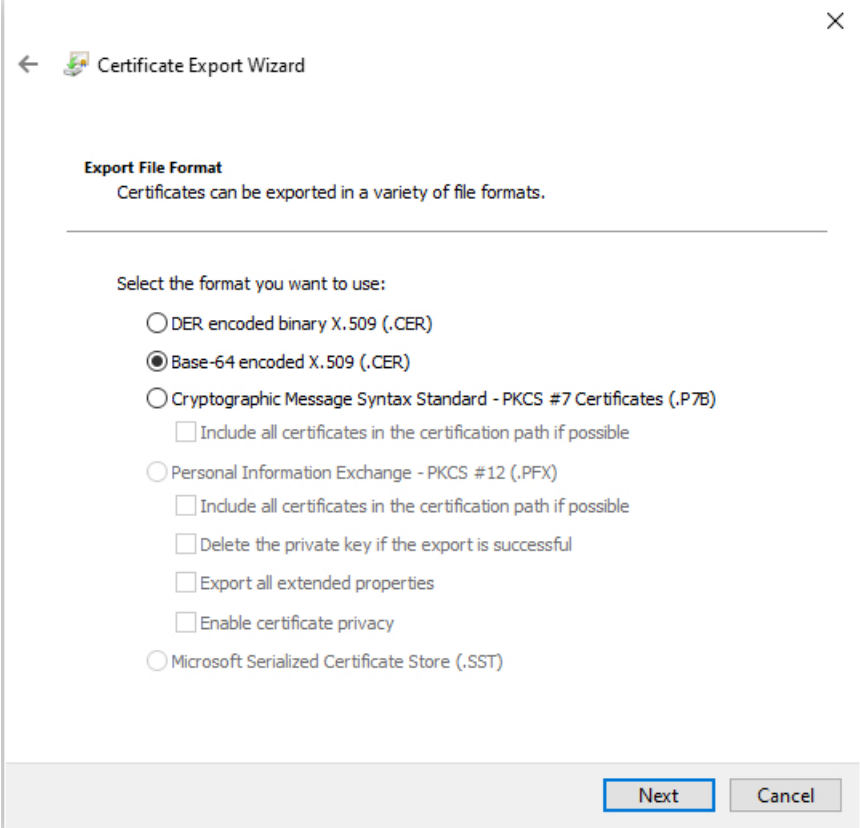

10. Follow the prompts to complete the export.
11. Open the certificate in a text editor.
12. Repeat the process for all certificates in the chain.

    You must paste each certificate in the text editor in order, first to last.
13. Repeat these tasks for vCenter or Firewall Management Center.

Get a Certificate Chain—Windows Firefox

Use the following procedure to get a certificate chain for the Firefox browser on either Windows or Mac OS.

1. Log in to vCenter or Firewall Management Center. using Firefox.
2. Click the lock to the left of the host name.
3. Click the right arrow (**Show connection details**). The following figure shows an example.

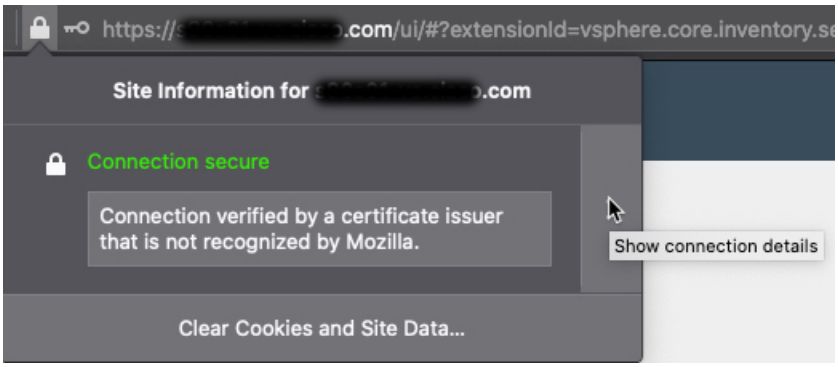

4. Click **More Information**.
5. Click **View Certificate**.
6. If the resulting dialog box has tab pages, click the tab page corresponding to the top-level CA.
7. Scroll to the Miscellaneous section.
8. Click **PEM (chain)** in the Download row. The following figure shows an example.

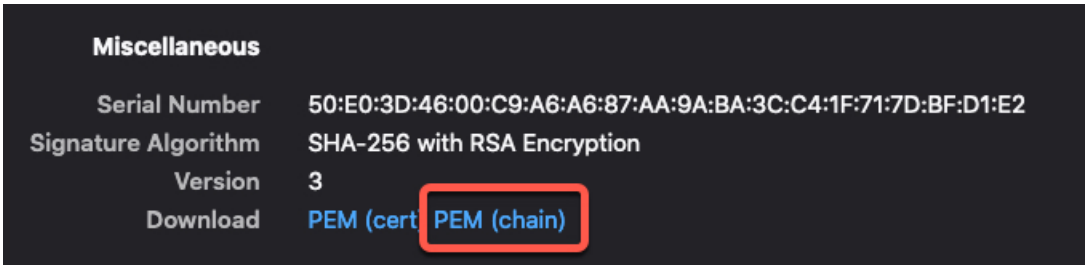

9. Save the file.
10. Repeat these tasks for vCenter or Firewall Management Center.

Security Requirements

To safeguard the Cisco Secure Dynamic Attributes Connector, you should install it on a protected internal network. Although the dynamic attributes connector is configured to have only the necessary services and ports available, you must make sure that attacks cannot reach it.

If the dynamic attributes connector and the Secure Firewall Management Center reside on the same network, you can connect the Secure Firewall Management Center to the same protected internal network as the dynamic attributes connector.

Regardless of how you deploy your appliances, inter-system communication is encrypted. However, you must still take steps to ensure that communications between appliances cannot be interrupted, blocked, or tampered with; for example, with a distributed denial of service (DDoS) or man-in-the-middle attack.

Internet Access Requirements

By default, the dynamic attributes connector is configured to communicate with the Firepower System over the internet using HTTPS on port 443/tcp (HTTPS). If you do not want the dynamic attributes connector to have direct access to the internet, you can configure a proxy server.

The following information informs you of the URLs the dynamic attributes connector use to communicate with the Secure Firewall Management Center and with external servers.

Table 99: Cisco Secure Dynamic Attributes Connector access requirements

URL	Reason
https://*fmc-ip***/api/fmc_platform/v1/auth/generatetoken**	Authentication
https://*fmc-ip***/api/fmc_config/v1/domain/***domain-id***/object/dynamicobjects**	GET and POST dynamic objects
https://*fmc-ip***/api/fmc_config/ v1/domain/***domain-id***/object/dynamicobjects/***object-id***/mappings?action=add**	Add mappings
https://*fmc-ip***/api/fmc_config/v1/domain/***domain-id* **/object/dynamicobjects/***object-id***/mappings?action=remove**	Remove mappings

Table 100: Cisco Secure Dynamic Attributes Connector vCenter access requirements

URL	Reason
https://*vcenter-ip***/rest/com/vmware/cis/session**	Authentication
https://*vcenter-ip***/rest/vcenter/vm**	Get VM information
https://*nsx-ip***/api/v1/fabric/virtual-machines/***vm-id*	Get NSX-T tag associated with the virtual machine

Migration from DockerHub to Amazon ECR

Docker images for the Cisco Secure Dynamic Attributes Connector are being migrated from Docker Hub to Amazon Elastic Container Registry (Amazon ECR).

To use the new field packages, you must allow access through your firewall or proxy to all of the following URLs:

- https://public.ecr.aws
- https://csdac-cosign.s3.us-west-1.amazonaws.com

Cisco Secure Dynamic Attributes Connector Azure access requirements

The dynamic attributes connector calls built-in SDK methods to get instance information. These methods internally call https://login.microsoft.com (for authentication) and https://management.azure.com (to get instance information).

History for the Cisco Secure Dynamic Attributes Connector

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
New connectors	7.6	7.6	AWS security groups, AWS service tags, and Cisco Cyber Vision These connectors can send an on-premises Secure Firewall Management Center dynamic objects as can Cisco Security Cloud Control. To receive dynamic objects from an on-premises dynamic attributes connector, version 3.0 of the on-premises dynamic attributes connector is required.
Cisco Secure Dynamic Attributes Connector	7.4.0	7.4.0	This feature is introduced. The Cisco Secure Dynamic Attributes Connector is now included in the Secure Firewall Management Center. You can use the dynamic attributes connector to get IP addresses from cloud-based platforms such as Microsoft Azure in access control rules without having to deploy to managed devices. More information: • The dynamic attributes connector included with this product: About the Cisco Secure Dynamic Attributes Connector, on page 1803 • The standalone dynamic attributes connector: Cisco Secure Dynamic Attributes Connector Configuration Guide New/modified screen: **Integration** > **Dynamic Attributes Connector**

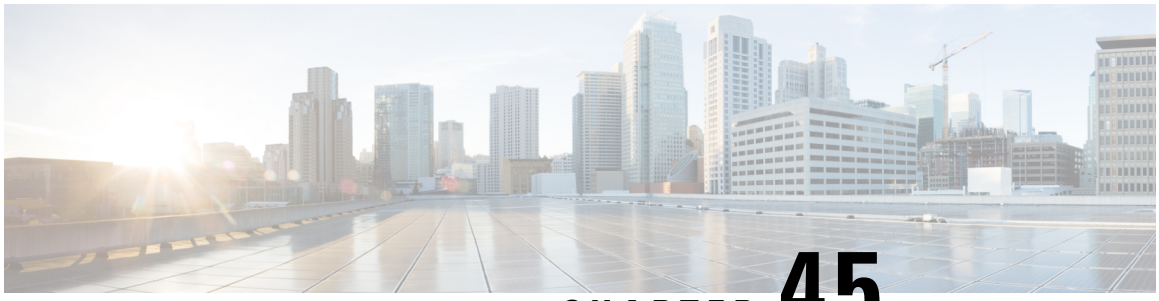

CHAPTER 45

URL Filtering

You can implement URL filtering using access control rules.

- URL Filtering Overview, on page 1857
- Best Practices for URL Filtering, on page 1859
- License Requirements for URL Filtering, on page 1864
- Requirements and Prerequisites for URL Filtering, on page 1864
- How to Configure URL Filtering with Category and Reputation, on page 1865
- Manual URL Filtering, on page 1871
- Configure HTTP Response Pages, on page 1873
- Configure URL Filtering Health Monitors, on page 1877
- Dispute URL Category and Reputation, on page 1877
- If the URL Category Set Changes, Take Action, on page 1878
- Troubleshoot URL Filtering, on page 1880
- History for URL Filtering, on page 1882

URL Filtering Overview

Use the URL filtering feature to control the websites that users on your network can access:

- Category and reputation-based URL filtering—With a URL Filtering license, you can control access to websites based on the URL's general classification (category) and risk level (reputation). This is the recommended option.
- Manual URL filtering—With any license, you can manually specify individual URLs, groups of URLs, and URL lists and feeds to achieve granular, custom control over web traffic. For more information, see Manual URL Filtering, on page 1871.

See also Security Intelligence, on page 1885, a similar but different feature for blocking malicious URLs, domains, and IP addresses.

The URL filtering feature is applicable only when you use a web browser to browse websites via the HTTP/HTTPS protocol. For other applications, you must use an FQDN-based access control rule instead of URL filtering. See Access Control Rules for more information.

About URL Filtering with Category and Reputation

With a URL Filtering license, you can control access to websites based on the category and reputation of requested URLs:

- Category—A general classification for the URL. For example, ebay.com belongs to the Auctions category, and monster.com belongs to the Job Search category.

  A URL can belong to more than one category.

- Reputation—How likely the URL is to be used for purposes that might be against your organization's security policy. Reputations range from Unknown risk (level 0) or Untrusted (level 1) to Trusted (level 5). Reputation is determined per complete Uniform Resource Identifier (URI), e.g. different pages on www.cisco.com can have different reputations.

Benefits of Category and Reputation-Based URL Filtering

URL categories and reputations help you quickly configure URL filtering. For example, you can use access control to block untrusted URLs in the Hacking category. Or, you can use QoS to rate limit traffic from sites in the Streaming Video category. There are also categories for types of threats, such as a Spyware and Adware category.

Using category and reputation data simplifies policy creation and administration. It grants you assurance that the system controls web traffic as expected. Because Cisco continually updates its threat intelligence with new URLs, as well as new categories and risks for existing URLs, the system uses up-to-date information to filter requested URLs. Sites that (for example) represent security threats, or that serve undesirable content, may appear and disappear faster than you can update and deploy new policies.

Some examples of how the system can adapt include:

- If an access control rule blocks all gaming sites, as new domains get registered and classified as Games, the system can block those sites automatically. Similarly, if a QoS rule rate limits all streaming video sites, the system can automatically limit traffic to new Streaming Video sites.
- If an access control rule blocks all malware sites and a shopping page gets infected with malware, the system can recategorize the URL from Shopping to Malware Sites and block that site.
- If an access control rule blocks untrusted social networking sites and somebody posts a link on their profile page that contains links to malicious payloads, the system can change the reputation of that page from Favorable to Untrusted and block it.

Limitations of category-based filtering in decryption policy Do Not Decrypt rules

You can optionally choose to include categories in your decryption policies. These categories, also referred to as *URL filtering*, are updated by the Cisco Talos intelligence group. Updates are based on machine learning and human analysis according to content that is retrievable from the website destination and sometimes from its hosting and registration information. Categorization is *not* based on the declared company vertical, intent, or security.

Note Don't confuse URL filtering with application detection, which relies on reading some of the packet from a website to determine more specifically what it is (for example, Facebook Message or Salesforce). For more information, see Best Practices for Configuring Application Control, on page 1735.

For more information, see Use Categories in URL Filtering, on page 1863.

URL Category and Reputation Descriptions

Category Descriptions

A description of each URL category is available from https://www.talosintelligence.com/categories.

Be sure to click **Threat Categories** to see those categories.

Reputation Level Descriptions

Go to Web reputation levels.

URL Filtering Data from the Cisco Cloud

Adding a URL Filtering license automatically enables the URL filtering feature. This allows traffic handling based on a website's general classification, or *category*, and risk level, or *reputation*.

By default, when users browse to an URL whose category and reputation is not in a local cache of previously accessed websites, the system submits it to the cloud for threat intelligence evaluation and adds the result to the cache.

Optionally, you can use a local URL dataset of categories and reputations, which can make web browsing faster. When you enable (or re-enable) URL filtering, the Firewall Management Center automatically queries Cisco for URL data and pushes the dataset to managed devices. Then, when users browse to an URL, the system checks the local dataset and the cache for category and reputation information before submitting it to the cloud for threat intelligence evaluation. To see your options for using the local dataset, including how to disable individual cloud lookups altogether, see URL Filtering Options, on page 1866.

Automatic updates of URL data is enabled by default; we strongly recommend you do not disable these updates.

The set of URL categories may change periodically. When you receive a change notification, review your URL filtering configurations to make sure traffic is handled as expected. For more information, see If the URL Category Set Changes, Take Action, on page 1878.

Best Practices for URL Filtering

Keep in mind the following guidelines and limitations for URL filtering:

Filter by Category and Reputation

Follow the instructions in How to Configure URL Filtering with Category and Reputation, on page 1865.

Configure Your Policy to Inspect Packets That Must Pass Before a URL Can Be Identified

The system cannot filter URLs before:

- A monitored connection is established between a client and server.
- The system identifies the DNS, HTTP or HTTPS application in the session.

- The system identifies the requested domain or URL (for encrypted sessions, from a non-encrypted domain name, the ClientHello message or the server certificate).

This identification should occur within 3 to 5 packets, or after the server certificate exchange in the TLS/SSL handshake if the traffic is encrypted.

Important! To ensure that your system examines these initial packets that would otherwise pass, see Inspection of Packets That Pass Before Traffic Is Identified, on page 2716 and subtopics.

If early traffic matches all other rule conditions but identification is incomplete, the system allows the packet to pass and the connection to be established (or the TLS/SSL handshake to complete). After the system completes its identification, the system applies the appropriate rule action to the remaining session traffic.

Block Threat Categories

Be sure that your policies specifically address Threat categories, which identify known malicious sites. Do this in addition to blocking sites with poor reputations.

For example, to protect your network from malicious sites, you must block all Threat categories. Additionally, Talos recommends that you block only sites with Poor category. You can block questionable reputations if you have an aggressive security posture, but this may result in a higher amount of false positives.

For specifics, see **Threat Categories** at the URL in URL Category and Reputation Descriptions, on page 1859.

URL Conditions and Rule Order

- Position URL rules after all other rules that *must* be hit.
- URLs can belong to more than one category. It is possible to want to allow one category of websites and block another—whether explicitly or by relying on the default action. In this case, make sure you create and order URL rules so you get the desired effect, depending on whether the allow or the block should take precedence.

For additional guidelines for rules, see the following topics: Best Practices for Access Control Rules, on page 1738.

Uncategorized or Reputationless URLs

When you build a URL rule, you first choose the category you want to match. If you explicitly choose **Uncategorized** URLs, you cannot further constrain by reputation.

Uncategorized URLs with Untrusted reputation are handled by the **Malicious Sites** category. If you want to block uncategorized sites with any other reputation level (such as Questionable), you must block all uncategorized sites.

After selecting a category and a reputation level, you can optionally select **Apply to unknown reputation**. For example, you can create a rule that applies to sites with Untrusted, Questionable, and unknown reputations.

You cannot manually assign categories and reputations to URLs, but in access control and QoS policies, you can manually block specific URLs. See Manual URL Filtering, on page 1871. See also Dispute URL Category and Reputation, on page 1877.

URL Filtering for Encrypted Web Traffic

When performing URL filtering on encrypted web traffic, the system:

- (If DNS filtering is enabled) Checks to see if the system has previously seen the originating domain or the domain is in the local reputation database, and if so, takes action based on the reputation and category of the domain. Otherwise, the system processes the traffic based on your configurations for encrypted traffic, even if **Retry URL cache miss lookup** is enabled in the access control policy's advanced settings.
- Disregards the encryption protocol; a rule matches both HTTPS and HTTP traffic if the rule has a URL condition but not an application condition that specifies the protocol.
- Does not use URL lists. You must use URL objects and groups instead.
- Matches HTTPS traffic based on the subject common name in the public key certificate used to encrypt the traffic, and also evaluates the reputation of any other URLs presented at any time during the transaction, including the post-decryption HTTP URL.
- Disregards subdomains within the subject common name.
- Does not display an HTTP response page for encrypted connections blocked by access control rules (or any other configuration); see Limitations to HTTP Response Pages, on page 1873.

URL Filtering and TLS Server Identity Discovery

The latest version of the Transport Layer Security (TLS) protocol 1.3, defined by RFC 8446, is the preferred protocol for many web servers to provide secure communications. Because the TLS 1.3 protocol encrypts the server's certificate for additional security, and the certificate is needed to match application and URL filtering criteria in access control rules, the Firepower System provides a way to extract the server certificate *without* decrypting the entire packet.

Access control policy advanced settings offer an **Early application detection and URL categorization** option for TLS Server Identity Discovery.

We strongly recommend enabling it for any traffic you want to match on application or URL criteria, especially if you want to perform deep inspection of that traffic. A decryption policy is not required because *traffic is not decrypted* in the process of extracting the server certificate.

Note

- TLS server identity discovery *cannot* be used with any of the following:
  - STARTTLS traffic
  - The HTTP CONNECT method
  - Traffic that is already being decrypted by another device on the network
- Because the certificate is decrypted, TLS server identity discovery can reduce performance depending on the hardware platform.
- TLS server identity discovery is not supported in inline tap mode or passive mode deployments.
- Enabling TLS server identity discovery is not supported on any Secure Firewall Threat Defense Virtual deployed to AWS. If you have any such managed devices managed by the Secure Firewall Management Center, the connection event **PROBE_FLOW_DROP_BYPASS_PROXY** increments every time the device attempts to extract the server certificate.
- TLS Server Identity Discovery also operates on TLS 1.2 sessions.

For more information, see Access Control Policy Advanced Settings, on page 1756.

HTTP/2

The system can extract HTTP/2 URLs from TLS certificates, but not from a payload.

Manual URL Filtering

- Specify URLs using a custom Security Intelligence list or feed object. Do not use a URL object or directly enter a URL into the rule. For details, see Manual URL Filtering Options, on page 1872.
- If you manually filter specific URLs using URL objects or by entering URLs directly into the rule, carefully consider other traffic that might be affected. To determine whether network traffic matches a URL condition, the system performs a simple substring match. If the requested URL matches any part of the string, the URLs are considered to match.
- If you use manual URL filtering to create exceptions to other rules, position the specific rule with the exceptions above the general rule that would otherwise apply.

Search Query Parameters in URLs

The system does not use search query parameters in the URL to match URL conditions. For example, consider a scenario where you block all shopping traffic. In that case, using a web search to search for amazon.com is not blocked, but browsing to amazon.com is.

URL Filtering in High Availability Deployments

For guidelines for URL filtering with Firepower Management Centers in high availability, see *URL Filtering and Security Intelligence* in the Cisco Secure Firewall Management Center Administration Guide .

Memory Limitations for Selected Device Models

- Device models with less memory store less URL data locally, and the system may therefore check the cloud more frequently to determine category and reputation for sites that are not in the local database.

  Lower-memory devices include:

  - Firepower 1010
  - Firewall Threat Defense Virtual with 8 GB of RAM

Filtering HTTPS Traffic

To filter encrypted traffic, the system determines the requested URL based on information passed during the TLS/SSL handshake: the subject common name in the public key certificate used to encrypt the traffic.

HTTPS filtering, unlike HTTP filtering, disregards subdomains within the subject common name. Do not include subdomain information when manually filtering HTTPS URLs in access control or QoS policies. For example, use example.com rather than www.example.com.

Tip In an decryption policies, you can handle and decrypt traffic to specific URLs by defining a distinguished name decryption policy rule condition. The common name attribute in a certificate's subject distinguished name contains the site's URL. Decrypting HTTPS traffic allows access control rules to evaluate the decrypted session, which improves URL filtering.

Controlling Traffic by Encryption Protocol

The system disregards the encryption protocol (HTTP vs HTTPS) when performing URL filtering in access control or QoS policies. This occurs for both manual and reputation-based URL conditions. In other words, URL filtering treats traffic to the following websites identically:

- http://example.com/
- https://example.com/

To configure a rule that matches only HTTP or HTTPS traffic, add an application condition to the rule. For example, you could allow HTTPS access to a site while disallowing HTTP access by constructing two access control rules, each with an application and URL condition.

The first rule allows HTTPS traffic to the website:

Action: Allow
Application: HTTPS
URL: example.com

The second rule blocks HTTP access to the same website:

Action: Block
Application: HTTP
URL: example.com

Use Categories in URL Filtering

Limitations of categories in Do Not Decrypt rules

You can optionally choose to include categories in your decryption policies. These categories, also referred to as *URL filtering*, are updated by the Cisco Talos intelligence group. Updates are based on machine learning and human analysis according to content that is retrievable from the website destination and sometimes from its hosting and registration information. Categorization is *not* based on the declared company vertical, intent, or security. While we strive to continuously update and improve URL filtering categories, it is not an exact science. Some websites are not categorized at all and it's possible some websites might be improperly categorized.

Avoid overusing categories in do not decrypt rules to avoid decrypting traffic without a reason; for example, the Health and Medicine category includes the WebMD website, which does not threaten patient privacy.

Following is a sample decryption policy that can prevent decryption for websites in the Health and Medicine categories but allows decryption for WebMD and everything else. General information about decryption rules can be found in Guidelines for Using TLS/SSL Decryption, on page 2304.

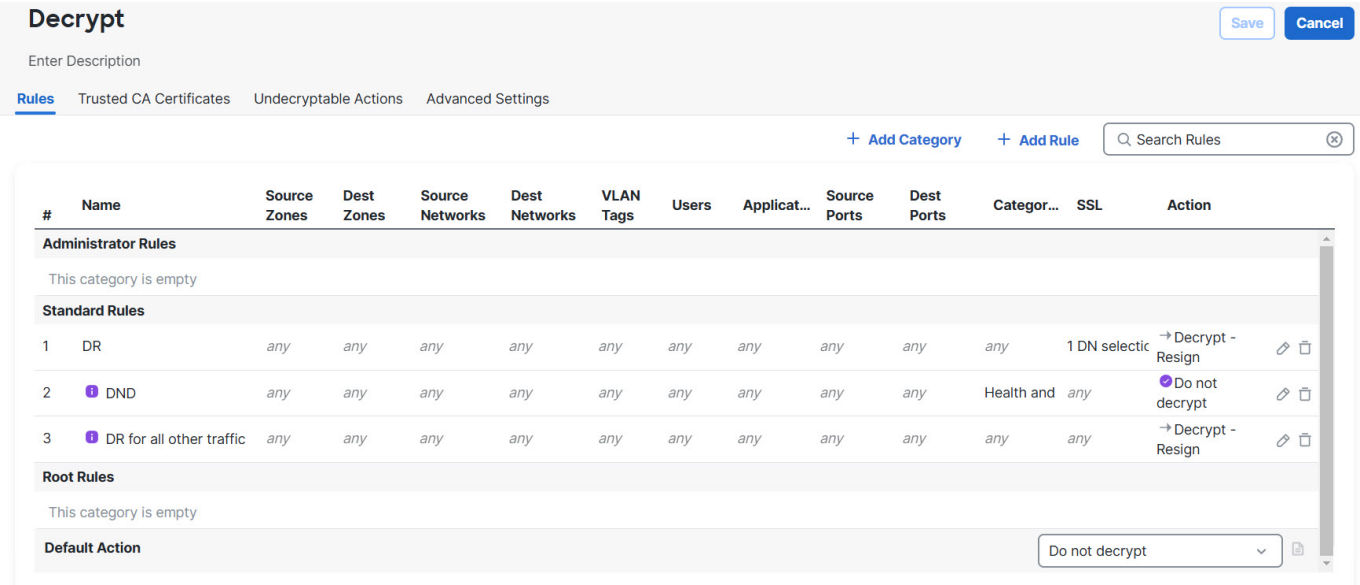

Note Don't confuse URL filtering with application detection, which relies on reading some of the packet from a website to determine more specifically what it is (for example, Facebook Message or Salesforce). For more information, see Best Practices for Configuring Application Control, on page 1735.

License Requirements for URL Filtering

Threat Defense License

- Category and reputation filtering—URL Filtering
- Manual filtering—No additional license.

URL Filtering Licenses for Threat Defense Devices

See *URL Licenses* in the *Licenses* chapter in the Cisco Secure Firewall Management Center Administration Guide.

Requirements and Prerequisites for URL Filtering

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

How to Configure URL Filtering with Category and Reputation

	Do This	More Information
Step 1	Ensure that you have the correct licenses.	Assign the URL Filtering license to each managed device that will filter URLs.
Step 2	Ensure that your Firewall Management Center can communicate with the cloud to obtain URL filtering data.	*Internet Access Requirements* and *Communication Port Requirements* in the Cisco Secure Firewall Management Center Administration Guide.
Step 3	Understand limitations and guidelines and take any necessary actions.	Best Practices for URL Filtering, on page 1859
Step 4	Enable the URL Filtering feature.	Enable URL Filtering Using Category and Reputation, on page 1866
Step 5	Configure rules to filter URLs by category and reputation.	Configuring URL Conditions, on page 1868 For the best protection against malicious sites, you must block sites by reputation AND block URLs in all Threat categories. (Optional) Supplement or Selectively Override Category and Reputation-Based URL Filtering, on page 1872
Step 6	(Optional) Allow users to bypass a website block by clicking through a warning page.	Configure HTTP Response Pages, on page 1873
Step 7	Order your rules so that traffic hits key rules first.	URL Rule Order, on page 1741

	Do This	More Information
Step 8	(Optional) Modify advanced options related to URL filtering.	Generally, use the defaults unless you have a specific reason to change them. For information about advanced options, including the following, see Access Control Policy Advanced Settings, on page 1756. • **Maximum URL characters to store in connection events** • **Allow an Interactive Block to bypass blocking for (seconds)** • **Retry URL cache miss lookup** • **Enable reputation enforcement on DNS traffic**
Step 9	Deploy your changes.	Deploy Configuration Changes, on page 259
Step 10	Ensure that your system receives future URL data updates as expected	Configure URL Filtering Health Monitors, on page 1877
Step 11	Be sure you have enabled other features that protect your network from malicious sites	See Security Intelligence, on page 1885.

Enable URL Filtering Using Category and Reputation

You must be an Admin user to perform this task.

Before you begin

Complete prerequisites described in How to Configure URL Filtering with Category and Reputation, on page 1865.

Procedure

Step 1 Choose **Integration** > **Other Integrations** > **Cloud Services**.

Step 2 Configure URL Filtering Options, on page 1866.

Step 3 Click **Save**.

URL Filtering Options

Adding a URL Filtering license automatically enables the URL filtering feature. This allows traffic handling based on a website's general classification, or *category*, and risk level, or *reputation*.

Although by default the system is configured to submit all URLs to the cloud for threat intelligence evaluation, using a local dataset of category and reputation data can make web browsing faster. When you enable (or re-enable) URL filtering, the Firewall Management Center automatically queries Cisco for URL data and pushes the dataset to managed devices. This process may take some time.

If you use SSL rules to handle encrypted traffic, also see Decryption Rule Guidelines and Limitations, on page 2304.

Enable Automatic Updates

If you **Enable Automatic Updates** (the default), the Firewall Management Center checks the cloud every 30 minutes for updates. If you need strict control over when the system contacts external resources, disable automatic updates and instead create a recurring task using the scheduler. See *Automated URL Filtering Updates Using a Scheduled Task* in the Cisco Secure Firewall Management Center Administration Guide .

Update Now

Click **Update Now** to perform a one-time, on-demand URL data update. You cannot start an on-demand update if an update is already in progress. Although daily updates tend to be small, if it has been more than five days since your last update, new URL data may take up to 20 minutes to download, depending on your bandwidth. Then, it may take up to 30 minutes to perform the update itself.

URL Query Source

You can choose how the system assigns a category and reputation to the URLs that your users browse to. You can choose:

- **Local Database Only**: Uses the local dataset only. Use this option if you do not want to submit your uncategorized URLs (category and reputation not in the local dataset) to Cisco, for example, for privacy reasons. However, note that connections to uncategorized URLs do *not* match rules with category or reputation-based URL conditions. You cannot assign categories or reputations to URLs manually.
- **Local Database and Cisco Cloud**: Uses the local dataset when possible, which can make web browsing faster. When users browse to an URL whose category and reputation is not in the local dataset or a cache of previously accessed websites, the system submits it to the cloud for threat intelligence evaluation and adds the result to the cache.
- **Cisco Cloud Only** (default): Does not use the local dataset. When users browse to an URL whose category and reputation is not in a local cache of previously accessed websites, the system submits it to the cloud for threat intelligence evaluation and adds the result to the cache. This option guarantees the most up-to-date category and reputation information.

  This option requires threat defense Version 7.3. If you enable this option, devices running earlier versions will use the **Local Database and Cisco Cloud** option.

Cached URLs Expire

This setting is not relevant if you chose **Local Database Only** as your URL query source.

Caching category and reputation data makes web browsing faster. By default, cached data for URLs never expires, for fastest performance.

To minimize instances of URLs matching on stale data, you can set URLs in the cache to expire. For greater accuracy and currency of threat data, choose a shorter expiration time. A cached URL refreshes *after* the first

time a user on the network accesses it after the specified time has passed. The first user does not see the refreshed result, but the next user who visits this URL does see the refreshed result.

Configuring URL Conditions

Protect your network by controlling access to sites based on URL category and reputation.

Before you begin

Attention As a prerequisite, ensure that you create at least a Monitor rule at the top of your access control policy, containing Category or Reputation parameters. This is essential to see ANY category or reputation data for ANY URLs that hit the particular access control policy.

If there is no rule in the access control policy with the category or reputation parameters configured, the **Connection Events** page in the management center shows no data for Category or Reputation for any URL traffic that hits the access control policy.

Procedure

Step 1 In the rule editor, click the following for URL conditions:

- Access control or QoS—Click **URLs**.
- SSL—Click **Category**.

Step 2 Find and choose the URL categories that you want to control:

In an access control or QoS rule, click **Category**.

For effective protection from malicious sites, you must block URLs in all Threat categories. Additionally, Talos recommends that you block only sites with Poor category. You can block questionable reputations if you have an aggressive security posture, but this may result in a higher amount of false positives. For a list of Threat categories, see URL Category and Reputation Descriptions, on page 1859.

Be sure to click the arrows at the bottom of the list to see all available categories.

Step 3 (Optional) Constrain URL categories by choosing a **Reputation**.

Note that if you explicitly match **Uncategorized** URLs, you cannot further constrain by reputation. Choosing a reputation level also includes other reputations either more or less severe than the level you choose, depending on the rule action:

- Includes less severe reputations—If the rule allows or trusts web traffic. For example, if you configure an access control rule to allow Favorable (level 4), it also automatically allows Trusted (level 5) sites.
- Includes more severe reputations—If the rule rate limits, decrypts, blocks, or monitors web traffic. For example, if you configure an access control rule to block Questionable sites (level 2), it also blocks Untrusted (level 1) sites.

If you change the rule action, the system automatically changes the reputation levels in URL conditions.

Optionally, select **Apply to unknown reputation**.

Step 4 Click **Add URL** or **Add to Rule**, or drag and drop.

Step 5 (Optional) To choose predefined URL objects, or URL lists and feeds in an access control or QoS rule, click **URL**, select the objects, and add them to the destination.

These objects implement manual URL filtering rather than category-based filtering.

Step 6 Save or continue editing the rule.

Example: URL Condition in an Access Control Rule

The following graphic shows the URL condition for an access control rule that blocks all malware sites, all untrusted sites, and all social networking sites with a reputation level of Neutral or worse.

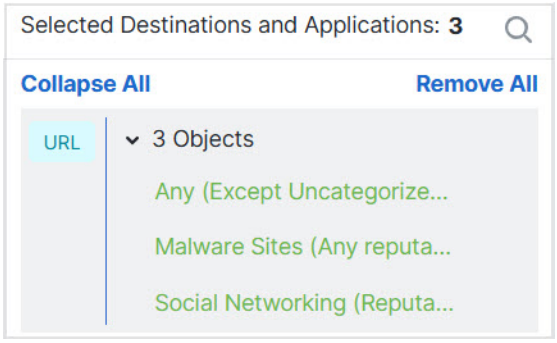

The following table summarizes how you build the condition.

Blocked URL	Category	Reputation
Malware sites, regardless of reputation	Malware Sites	Any
Any untrusted URL (level 1)	Any	1 - Untrusted
Social networking sites with a reputation level of Neutral or worse (levels 1 through 3)	Social Network	3 - Neutral

Rules with URL Conditions

The following table lists rules that support URL conditions, and the types of filtering that each rule type supports.

Rule Type	Supports Category and Reputation Filtering?	Supports Manual Filtering?
Access control	Yes	Yes
Decryption policy	Yes	No; use distinguished name conditions instead
QoS	Yes	Yes

To use URL filtering in a decryption policy that has **Do Not Decrypt** rule conditions, see Use Categories in URL Filtering, on page 1863.

URL Rule Order

For the most effective URL matching, place rules that include URL conditions before other rules, particularly if the URL rules are block rules and the other rules meet both of the following criteria:

- They include application conditions.
- The traffic to be inspected is encrypted.

If you configure exceptions to a rule, put the exception above the other rule.

DNS Filtering: Identify URL Reputation and Category During DNS Lookup

The **Enable reputation enforcement on DNS traffic** option is enabled by default on the **Advanced** tab of each new access control policy. This option slightly modifies URL filtering behavior and is applicable only when URL filtering is enabled and configured.

When this option is enabled:

- The system evaluates domain category and reputation early in URL transactions, when the browser looks up the domain name to get the IP address
- Category and reputation of encrypted traffic can often be determined without decryption

  If DNS filtering cannot determine the URL of encrypted traffic, that traffic is processed using your configurations for encrypted traffic.

Enable DNS Filtering to Identify URLs During Domain Lookup

DNS filtering is enabled by default in new access control policies. However, additional configurations may be required in order for this setting to take effect.

Before you begin

- URL filtering using category and reputation must be licensed, enabled, and configured.

  (DNS filtering does not use the following settings in the URLs tab: URL groups, URL objects, URL lists and feeds, and URLs entered into the "Enter URL" text box.)
- See limitations at DNS Filtering Limitations, on page 1871.

Procedure

Step 1 In your access control policy's advanced settings, select **Enable reputation enforcement on DNS traffic**.

Step 2 In the same policy, for each access control rule that has URL category and reputation blocking configured:

- Application conditions—If the application condition is anything other than **any** (or empty), add **DNS** to that list. Other DNS-related options are not relevant for this purpose.

- Port condition—If the port/protocol condition is anything other than **any** (or empty), add **DNS_over_TCP** and **DNS_over_UDP**.

Step 3 Save your changes.

What to do next

If you are done making changes: Deploy Configuration Changes, on page 259.

DNS Filtering Limitations

Traffic that matches rules having action **Block with reset**, **Interactive Block**, or **Interactive Block with reset** will be treated as if the rule action were **Block**.

End users trying to access a blocked URL will experience this as an unexplained inability to connect to their page; the connection will spin and then time out.

DNS Filtering and Events

Connection events generated by DNS filtering are logged using the following fields: DNS Query, URL Category, URL Reputation, and Destination Port. The DNS Query field holds the domain name; the URL field will be blank for DNS filtering matches. The Destination Port will be 53.

Also:

- When the access control rule action is **Allow** or **Trust**, two connection events will be generated for the same traffic, one for DNS filtering (with the **DNS Query** field populated) and one for URL filtering (with the **URL** field populated).
- The first time the system encounters a particular URL, you will see two events for that single session: One event showing uncategorized/reputationless for the DNS Query, and one event showing the actual category and reputation for the URL, which were retrieved during the DNS Query and applied to the session while processing using standard URL filtering.

Manual URL Filtering

In access control and QoS rules, you can supplement or selectively override category and reputation-based URL filtering by manually filtering individual URLs, groups of URLs, or URL lists and feeds.

Examples:

- You might use access control to block a category of websites that are not appropriate for your organization. However, if the category contains a website that is appropriate, and to which you want to provide access, you can create a manual Allow rule for that site and place it before the Block rule for the category.
- If you prefer not to use category-based URL filtering, you can manually permit or block a URL or a group of URLs by creating an Allow or Block rule within an access control policy. You can perform this type of URL filtering without a special license.

Manual URL filtering is not supported in SSL rules; instead, use distinguished name conditions.

Caution Depending on how you implement manual URL filtering, URL matching may not be what you intend. See Manual URL Filtering Options, on page 1872.

Manual URL Filtering Options

There are several ways to specify URLs for manual URL filtering:

Option	Description
(Best practice) Use custom Security Intelligence URL list or feed objects.	This is the recommended method for manual URL filtering. You can create a new list or feed, or choose an existing one in an access control or QoS rule. For more information, see Custom Security Intelligence Lists and Feeds, on page 1435 and subtopics.
Use URL objects, individually or as groups. URL objects are described at URL, on page 1444. Or Enter URLs directly into the access control rule. (The **Enter URL** option on the rule page in the web interface.)	If you do not include a path (that is, there are no / characters in the URL), the match is based on the server's hostname only. If you include one or more / character, the entire URL string is used for a substring match. Then, a URL is considered a match if any of the following are true: • The string is at the beginning of the URL. • The string follows a dot. • The string contains a dot in the beginning. • The string follows the :// characters. For example, ign.com matches ign.com or www.ign.com, but not versign.com. **Note** We recommend that you do not use manual URL filtering to block or allow individual web pages or parts of sites (that is, URL strings with / characters), as servers can be reorganized and pages moved to new paths. The **Enter URL** option does not support wildcards.

Supplement or Selectively Override Category and Reputation-Based URL Filtering

In access control or QoS rules, you can use Security Intelligence URL lists and feeds to supplement, or to specify exceptions to, your category and reputation-based URL filtering rules.

Important! If the list or feed you are configuring in this procedure contains exceptions to category- or reputation-based rules, put this rule above those rules in the rule order.

In SSL rules, use distinguished name conditions to configure parallel behavior.

Before you begin

- Configure URL filtering using category and reputation. See Configuring URL Conditions, on page 1868.
- Understand important best practices for manual URL filtering. See Best Practices for URL Filtering, on page 1859 and Manual URL Filtering Options, on page 1872.
- Configure one or more Security Intelligence objects (lists or feeds) containing the URLs that you want to use for manual filtering. See Custom Security Intelligence Lists and Feeds, on page 1435.

Procedure

Step 1 Navigate to the access control or QoS policy in which you will define the rule.

Step 2 Create or edit the rule in which you will add the new condition:

- If you are supplementing a category- or reputation-based URL filtering rule, edit the existing rule.
- If you are overriding or creating exceptions to a category- or reputation-based URL filtering rule, create a new rule.

Step 3 Select the list or feed you created as the destination URL criteria.

Step 4 Save the rule.

Configure HTTP Response Pages

As part of access control, you can configure an *HTTP response page* to display when the system blocks web requests, using either access control rules or the access control policy default action.

The response page displayed depends on how you block the session:

- **Block Response Page**: Overrides the default browser or server page that explains that the connection was denied.
- **Interactive Block Response Page**: Warns users, but also allows them to click a button (or refresh the page) to load the originally requested site. Users may have to refresh after bypassing the response page to load page elements that did not load.

If you do not choose a response page, the system blocks sessions without interaction or explanation.

Limitations to HTTP Response Pages

- The system displays a response page only for unencrypted or decrypted HTTP/HTTPS connections blocked (or interactively blocked) either by access control rules or by the access control policy default action. The system does not display a response page for connections that are blocked by any other policy or mechanism.
- The system cannot display a response page if the connection is reset (RST packet sent). If you enable response pages, the system prioritizes that configuration. Even if you choose **Block with reset** or **Interactive Block with reset** as the rule action, the system displays the response page and does not reset

matching web connections. To ensure that blocked web connections are reset, you must disable response pages.

Note that all non-web traffic that matches the rule *is* blocked with reset.

- The system does not display a response page for encrypted connections that are blocked by access control rules (or any other configuration). Access control rules evaluate encrypted connections if you did not configure an SSL policy, or your SSL policy passes encrypted traffic.

  For example, the system cannot decrypt HTTP/2 or SPDY sessions. If web traffic encrypted using one of these protocols reaches access control rule evaluation, the system does not display a response page if the session is blocked.

  However, the system does display a response page for connections that are decrypted by the SSL policy, then blocked (or interactively blocked) either by access control rules or by the access control policy default action. In these cases, the system encrypts the response page and sends it at the end of the reencrypted SSL stream.

- The system does not display a response page when web traffic is blocked because of a promoted access control rule (an early-placed blocking rule with only simple network conditions).
- If a URL is entered without specifying "http" or "https", and the browser initiates the connection on port 80, and the user clicks through a response page, and the connection is subsequently redirected to port 443, the user will not see a second interactive response page because the response to this URL is already cached.
- The system does not display a response page when web traffic is blocked before the system identifies the requested URL; see Best Practices for URL Filtering, on page 1859.
- The system does not display a response page if the block URL access control rule is configured after the allow application rule.

Requirements and Prerequisites for HTTP Response Pages

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Choosing HTTP Response Pages

Reliable display of HTTP response pages depends on your network configuration, traffic loads, and size of the page. Smaller pages are more likely to display successfully.

Procedure

Step 1 In the access control policy editor, select **HTTP Responses** from the **More** drop-down arrow at the end of the packet flow line.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 Choose the **Block Response Page** and **Interactive Block Response Page**:

- System-provided—Displays a generic response. Click **View** ( ) to view the code for this page.
- Custom—Create a custom response page. A pop-up window appears, prepopulated with system-provided code that you can replace or modify by clicking **Edit** ( ). A counter shows how many characters you have used.
- None—Disables the response page and blocks sessions without interaction or explanation. To quickly disable interactive blocking for the whole access control policy, choose this option.

Step 3 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Configure Interactive Blocking with HTTP Response Pages

When you configure interactive blocking, users can load an originally requested site after reading a warning. Users may have to refresh after bypassing the response page to load page elements that did not load.

Tip To quickly disable interactive blocking for the whole access control policy, display neither the system-provided page nor a custom page. The system then blocks all connections without interaction.

If a user does not bypass an interactive block, matching traffic is denied without further inspection. If a user bypasses an interactive block, the access control rule allows the traffic, although the traffic may still be subject to deep inspection and blocking.

By default, a user bypass is in effect for 10 minutes (600 seconds) without displaying the warning page on subsequent visits. You can set the duration to as long as a year, or you can force the user to bypass the block every time. This limit applies to every Interactive Block rule in the policy. You cannot set the limit per rule.

Logging options for interactively blocked traffic are identical to those in allowed traffic, but if a user does not bypass the interactive block, the system can log only beginning-of-connection events. When the system initially warns the user, it marks any logged beginning-of-connection event with the `Interactive Block`

or `Interactive Block with reset` action. If the user bypasses the block, additional connection events logged for the session have an action of `Allow`.

Configuring Interactive Blocking

The following procedure explains how to allow users to bypass URL filtering rules.

Procedure

Step 1 As part of access control, configure an access control rule that matches web traffic; see Create and Edit Access Control Rules, on page 1783:

- Action—Set the rule action to **Interactive Block** or **Interactive Block with reset**; see Access Control Rule Interactive Blocking Actions, on page 1780.
- Conditions—Use URL conditions to specify the web traffic to interactively block; see URL Conditions (URL Filtering).
- Logging—Assume users will bypass the block and choose logging options accordingly; see *Logging for Allowed Connections* in the Cisco Secure Firewall Management Center Administration Guide .
- Inspection—Assume users will bypass the block and choose deep inspection options accordingly; see Access Control Overview, on page 1723.

Step 2 (Optional) In the access control policy **HTTP Responses**, choose a custom interactive-block HTTP response page; see Choosing HTTP Response Pages, on page 1875.

Step 3 (Optional) In access control policy **Advanced** settings, change the user bypass timeout; see Setting the User Bypass Timeout for a Blocked Website, on page 1876.

After a user bypasses a block, the system allows the user to browse to that page without warning until the timeout period elapses.

Step 4 Save the access control policy.

Step 5 Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Setting the User Bypass Timeout for a Blocked Website

The following procedure explains how to set the time allowed for browsing after the user bypasses a URL filtering block. After the timeout expires, the user must bypass the block again.

Procedure

Step 1 Click **Policies** > **Access Control heading** > **Access Control** and edit the policy.

Step 2 Select **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 3 Click **Edit** (✎) next to General Settings.

If **View** (👁) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 4 In the **Allow an Interactive Block to bypass blocking for (seconds)** field, enter the number of seconds that must elapse before the user bypass expires.

Setting this value to 0 means the interactive block response is displayed once and the user bypass never expires.

Step 5 Click **OK**.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Configure URL Filtering Health Monitors

The following health policies alert if the system has problems obtaining or updating URL category and reputation data.

- URL Filtering Monitor
- Threat Data Updates on Device

To ensure that these are configured the way you want them, see *Health Modules* and *Configuring Health Monitoring* in the Cisco Secure Firewall Management Center Administration Guide.

Dispute URL Category and Reputation

If you disagree with a category or reputation assigned by Talos, you can submit a request for re-evaluation.

Before you begin

You will need your Cisco account credentials.

Procedure

Step 1 In the Firewall Management Center web interface, do one of the following:

Location of Dispute Option	Path to Dispute Option
Cloud Services configuration page	a. Navigate to the **Integration** > **Other Integrations** > **Cloud Services** page. b. Select **Dispute URL categories and reputations**.

Location of Dispute Option	Path to Dispute Option
Manual URL Lookup page	a. Navigate to the manual URL Lookup page: **Analysis** > **Advanced** > **URL**. b. Look up the URL in question. c. To see **Dispute** at the end of the table row, hover over the relevant entry in the list of results, then click dispute.
URL Connection Event	a. Navigate to any page under the **Analysis > Connections** menu that has a table that includes URLs. b. Right-click an item in the **URL Category** or **URL Reputation** column (show hidden columns if needed) and select an option.

The Talos web site opens in a separate browser window.

Step 2 Sign in to the Talos site with your Cisco credentials.

Step 3 Review the information and follow the instructions on the Talos page.

Step 4 Look for information on the Talos site about how submitted disputes are handled and what response to expect, if any.

The dispute process is independent of Firepower products.

If the URL Category Set Changes, Take Action

The set of URL Filtering categories may occasionally change, in order to accommodate new web trends and evolving usage patterns.

These changes affect both policies and events.

Shortly before URL category changes are scheduled to occur, and after they occur, you will see alerts in the list of rules in any access control, SSL, and QoS policy that is affected by the changes, and on URL or Category in rules that you edit.

You should take action when you see these alerts.

Note Updates to the URL category set as described in this topic are distinct from the changes that simply add new URLs to existing categories or re-classify misclassified URLs. This topic does not apply to category changes for individual URLs.

Procedure

Step 1 If you see an alert beside a rule in an access control policy, hover over the alert to see details.

Step 2 If the alert mentions changes to URL categories, edit the rule to see further details.

Step 3 Hover over the URL or Category in the rule dialog to see general information about the type of changes.

Step 4 If you see an alert beside a category, click the alert to view details.

Step 5 If you see a "More information" link in the description of a change, click it to view information about the category on the Talos web site.

Alternately, see a list and descriptions all categories at the link in URL Category and Reputation Descriptions, on page 1859.

Step 6 Depending on the type of change, take appropriate action:

Type of Category Change	What The System Will Do	What You Should Do
Existing category will soon be deprecated	Nothing yet. You have a few weeks to change affected rules. If you do not take action in that time, the system eventually will not be able to redeploy the policy.	Remove this category from all rules that include it. If there is a similar new category, consider using that category instead.
New category is added	By default, the system does not use newly added categories.	Consider creating new rules for the new category.
Existing category is deleted	The category will appear in the rule in strikethrough text (that is, with a line through the category name.)	You must delete the obsolete category from the rule before you can deploy the policy.

Step 7 Check your SSL rules (Category) for these changes and take action as needed.

Step 8 Check your QoS rules (URL) for these changes and take action as needed.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 259.

URL Category and Reputation Changes: Effect on Events

- When URL categories change, events that the system processed before the category change will be associated with their original category names and will be labeled with **Legacy**. Events that the system processed after the category change will be associated with the new categories.

  Older, legacy events will age out of the system over time.

- If a URL does not have a reputation at the time it was processed, the URL Reputation column in the event viewer will be empty.

Troubleshoot URL Filtering

Expected URL Category is Missing from the Categories List

The URL filtering feature uses a different set of categories than the Security Intelligence feature; the category that you expect to see may be a Security Intelligence category. To see those categories, look at the **URLs** tab on the **Security Intelligence** tab in an access control policy.

Initial Packets Are Passing Uninspected

See Inspection of Packets That Pass Before Traffic Is Identified, on page 2716 and subtopics.

See also DNS Filtering: Identify URL Reputation and Category During DNS Lookup , on page 1870.

Health alert: "URL Filtering registration failure"

Verify that your Firewall Management Center and any proxies can connect to the Cisco cloud. You may need information about URL Filtering and URL categories in the following topics: *Internet Access Requirements* and *Communication Port Requirements* in the Cisco Secure Firewall Management Center Administration Guide.

How can I find the category and reputation of a particular URL?

Do a manual lookup. See *Finding URL Category and Reputation* in the Cisco Secure Firewall Management Center Administration Guide.

Error when attempting a manual lookup: "Cloud Lookup Failure for <URL>"

Make sure the feature is properly enabled. See the prerequisites in *Finding URL Category and Reputation* in the Cisco Secure Firewall Management Center Administration Guide.

URL appears to be incorrectly handled based on its URL category and reputation

Problem: The system does not handle the URL correctly based on its URL category and reputation.

Solutions:

- Verify that the URL category and reputation associated with the URL are what you think they are. See *Finding URL Category and Reputation* in the Cisco Secure Firewall Management Center Administration Guide.
- The following issues may be addressed by settings described in URL Filtering Options, on page 1866, accessible using Enable URL Filtering Using Category and Reputation, on page 1866.
  - The URL cache may hold stale information. See information about the **Cached URLs Expire** setting in URL Filtering Options, on page 1866.
  - The local data set may not be updated with current information from the cloud. See information about the **Enable Automatic Updates** setting in URL Filtering Options, on page 1866.
  - The system may be configured to *not* check the cloud for current data. See information about the **Query Cisco cloud for unknown URLs** setting in URL Filtering Options, on page 1866.

- Your access control policy may be configured to pass traffic to the URL without checking the cloud. See information about the **Retry URL cache miss lookup** setting in Access Control Policy Advanced Settings, on page 1756.
- See also Best Practices for URL Filtering, on page 1859.
- If the URL is processed using an SSL rule, see Decryption Rule Guidelines and Limitations, on page 2304 and Decryption Policy Rule Order
- Verify that the URL is being handled using the access control rule that you think it is being handled by, and that the rule does what you think it does. Consider rule order.
- Verify that the local URL category and reputation database on the Firewall Management Center is successfully being updated from the cloud and that managed devices are successfully being updated from the Firewall Management Center.

  Status of these processes are reported in the Health Monitor, in the **URL Filtering Monitor** module and the **Threat Data Updates on Devices** module. For details, see *Health* in the Cisco Secure Firewall Management Center Administration Guide .

If you want to immediately update the local URL category and reputation database, go to **Integration** > **Other Integrations** > **Cloud Services**, then click **Update Now** in the **URL Filtering** area. For more information, see URL Filtering Options, on page 1866.

A URL category or reputation is not correct

For access control or QoS rules: Use manual filtering, paying careful attention to rule order. See Manual URL Filtering, on page 1871 and Configuring URL Conditions, on page 1868.

For SSL rules: Manual filtering is not supported. Instead, use distinguished name conditions.

See also Dispute URL Category and Reputation, on page 1877.

Web pages are slow to load

There is a tradeoff between security and performance. Some options:

- Consider modifying the **Cached URLs Expire** setting. Choose **Integration** > **Other Integrations** > **Cloud Services**. For information, see URL Filtering Options, on page 1866.
- Consider deselecting the **Retry URL cache miss lookup** setting in Access Control Policy Advanced Settings, on page 1756.

Events Do Not Include URL Category and Reputation

- Make sure you have included applicable URL rules in an access control policy, the rules are active, and the policies have been deployed to the relevant devices.
- URL category and reputation do not appear in an event if the connection is processed before it matches a URL rule.
- The rule that handles the connection must be configured for URL category and reputation.
- Even if you have configured URL categories in the Categories tab in an SSL rule, you must also configure the URLs tab in a rule in your access control policy.

DNS Filtering is not working

Make sure you have completed all prerequisites and steps in Enable DNS Filtering to Identify URLs During Domain Lookup , on page 1870.

An End User Tries to Access a Blocked URL and the Page Just Spins and Times Out

When DNS Filtering is enabled and end users access a URL that is blocked, the page will spin but not load. End users are not notified that the page is blocked. This is currently a limitation when DNS filtering is enabled.

See DNS Filtering Limitations, on page 1871.

Events Include URL Category and Reputation but URL Field is Blank

If the DNS Query field is populated and the URL field is empty, this is expected when the DNS filtering feature is enabled.

See DNS Filtering and Events, on page 1871.

Multiple Events are Generated for a Single Transaction

A single web transaction sometimes generates two connection events, one for DNS filtering and one for URL filtering. This is expected when DNS filtering is enabled and:

- the access control rule action for the traffic is Allow or Trust.
- the system encounters a URL for the first time.

See DNS Filtering and Events, on page 1871.

History for URL Filtering

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
New URL category	New in release 7.0 timeframe, applies to all releases	Any	New URL category: Private IP address For details, see Talosintelligence.com
DNS filtering	7.0 6.7 (Beta)	Any	A new option in the advanced settings for each access control policy allows earlier filtering of web traffic by category and reputation. This feature is enabled by default on new installations. Supported Platforms: Firewall Management Center and managed devices at any supported version.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Ability to specify handling for sites with unknown reputation	6.7	Any	You can now specify handling for URLs with unknown reputation. Modified screens: URL rules in access control policies and QoS policies, and category rules in SSL policies, include a new checkbox for this purpose below the reputation selection area. Supported Platforms: All
New and changed URL categories New names for reputation levels	6.5	Any	The following changes apply to URL rules in access control and QoS policies and to Category rules in SSL policies: The set of URL categories has changed. There are now two "pages" of categories from which to select when you create a URL rule. The name associated with each reputation level has changed. For descriptions of the new categories and reputation names, see URL Category and Reputation Descriptions, on page 1859. For complete details specific to upgrades, see also the Release Notes and upgrade instructions for version 6.5. If there are future category set changes, your rules will display icons to alert you. Modified screens: URL rules in access control policies, SSL policies, and QoS policies; event data related to URL categories. Supported Platforms: Firewall Management Center and devices running release 6.5.
Minor change to classic device licensing	6.5	Any	For devices that use classic licenses, URL filtering will not be enabled until the device is registered to the Firewall Management Center and a URL Filtering license is assigned to the device. Supported Platforms: NGIPSv and ASA with FirePOWER Services devices.
Addresses for retrieving URL data from the Cisco cloud have changed	6.5	Any	See the URL Filtering row in *Internet Access Requirements* in the Cisco Secure Firewall Management Center Administration Guide .

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Opportunity to dispute an assigned URL Category	6.5	Any	If you disagree with the category that the system assigns to a URL, you can submit a request to change the category. New/Modified screens: • New menu option when right-clicking a URL category or reputation in tables of connection events under the **Analysis** menu. • New button on the URL Lookups page (**Analysis** > **Advanced** > **URL**). (Hover your pointer over the URL to display the button.) • New option on the **System** > **Integration** > **Cloud Services** page Supported platforms: All
The **Cisco CSI** tab is renamed to **Cloud Services**	6.4	Any	Modified screens and navigation:**System** > **Integration** > **Cisco CSI** is now **System** > **Integration** > **Cloud Services** Supported platforms: Firewall Management Center
Moved URL Filtering information from various locations to this new URL Filtering chapter	6.3	Any	Moved information about configuring cloud communications for URL Filtering to the new URL Filtering chapter. Moved certain other URL Filtering information from other locations to this chapter. Made related changes to the structure of the Cisco CSI topics in the chapter.
New option: Cached URLs Expire	6.3	Any	Use this new control to balance performance with freshness of URL category and reputation data in order to minimize instances of URLs matching on stale data. Modified screens: **System** > **Integration** > **Cisco CSI**. Supported Platforms: All.
Changed menu path	6.3	Any	The path to the manual URL Lookup page has changed from **Analysis** > **Lookup** > **URL** to **Analysis** > **Advanced** > **URL**.

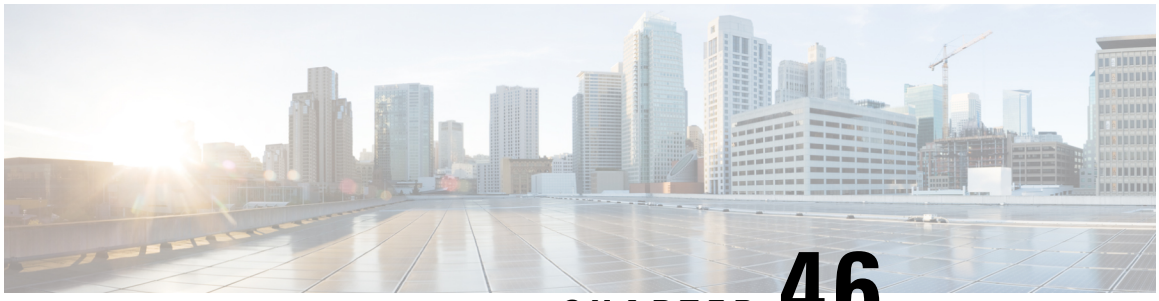

CHAPTER 46

Security Intelligence

The following topics provide an overview of Security Intelligence, including use of lists for blocking and allowing traffic and basic configuration.

- About Security Intelligence, on page 1885
- Best Practices for Security Intelligence, on page 1886
- License Requirements for Security Intelligence, on page 1886
- Requirements and Prerequisites for Security Intelligence, on page 1887
- Security Intelligence Sources, on page 1887
- Configure Security Intelligence, on page 1888
- Security Intelligence Monitoring, on page 1894
- Override Security Intelligence Blocking, on page 1895
- Troubleshooting Security Intelligence, on page 1895
- History for Security Intelligence Block Listing, on page 1896

About Security Intelligence

As an early line of defense against malicious internet content, Security Intelligence uses reputation intelligence to quickly block connections to or from IP addresses, URLs, and domain names. This is called *Security Intelligence block listing*.

Security Intelligence is an early phase of access control, before the system performs more resource-intensive evaluation. Using a Block list improves performance by quickly excluding traffic that does not require inspection.

Note You cannot use a Block list to block fastpathed traffic. Prefilter evaluation occurs before Security Intelligence filtering. Fastpathed traffic bypasses all further evaluation, including Security Intelligence.

Although you can configure custom Block lists, Cisco provides access to regularly updated intelligence feeds. Sites representing security threats such as malware, spam, botnets, and phishing appear and disappear faster than you can update and deploy custom configurations.

You can refine Security Intelligence Block listing with Do Not Block lists and monitor-only Block lists. These mechanisms exempt traffic from being blocked by a Block list, but do **not** automatically trust or fastpath matching traffic. Traffic added to a Do Not Block list or monitored at the Security Intelligence stage is intentionally subject to further analysis with the rest of access control.

Note The Security Intelligence configuration that includes detection signatures or malware URLs might lead to false threat alerts when the Firewall Threat Defense troubleshoot files are scanned.

Related Topics

Security Intelligence, on page 1429

Best Practices for Security Intelligence

- Configure your access control policies to block threats detected by Cisco-provided Security Intelligence feeds. See Configuration Example: Security Intelligence Blocking, on page 1893.
- If you want to supplement the Cisco-provided Security Intelligence feeds with custom threat data, or manually block emerging threats:
  - For IP addresses, use custom Security Intelligence lists and feeds, or Network objects or groups. To create these, see Security Intelligence, on page 1429 and Network, on page 1401, and their subtopics. To use them for Security Intelligence, see Configure Security Intelligence, on page 1888. Network objects used in Security Intelligence policy require an IPS license.
  - For URLs and domains, use custom Security Intelligence lists and feeds, *not* objects or groups. See details at Manual URL Filtering Options, on page 1872.
  - You can also add entries to a Block list from events. See Global and Domain Security Intelligence Lists, on page 1431.
- To test new feeds, or for passive deployments, set the action from block to monitor only. See Security Intelligence Monitoring, on page 1894.
- If you need to exclude specific sites or addresses from Security Intelligence blocking, see Override Security Intelligence Blocking, on page 1895.
- If your Secure Firewall deployment is integrated with Cisco security cloud and Cisco XDR and you use custom Security Intelligence lists and feeds, ensure that you update security services exchange with these lists and feeds. For details, see instructions for configuring auto-promotion of events in the Security Services Exchange online help. For general information about this integration, see *Integrate with Cisco Security Cloud* in the Cisco Secure Firewall Management Center Administration Guide.
- System-provided Security Intelligence categories may change over time and without notification; you should plan to check periodically for changes, and modify your policies accordingly.
- You should also configure URL filtering, a separate feature with separate licensing requirements, for further protection against malicious sites. See URL Filtering, on page 1857.

License Requirements for Security Intelligence

Threat Defense License

IPS

Requirements and Prerequisites for Security Intelligence

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Important You must apply the Network Discovery policy on the device for a successful application of the SI policy.

Security Intelligence Sources

- System-provided feeds

  Cisco provides access to regularly updated intelligence feeds for domains, URLs and IP addresses. For more information, see Security Intelligence, on page 1429.

  If you see a feed with "TID" in the name, this feed is *not* used by Security Intelligence. Instead, this feed is used by the feature described in Secure Firewall Threat Intelligence Director, on page 2735.

- Third-party feeds

  Optionally, supplement Cisco-provided feeds with third-party reputation feeds, which are dynamic lists that the Secure Firewall Management Center downloads from the internet on a regular basis. See Custom Security Intelligence Feeds, on page 1437.

- Custom Block lists or feeds (or objects or groups)

  Block specific IP addresses, URLs, or domain names using a manually-created list or feed (for IP addresses, you can also use network objects or groups.)

  For example, if you become aware of malicious sites or addresses that are not yet blocked by a feed, add these sites to a custom Security Intelligence list and add this custom list to the Block list in the Security Intelligence tab of your access control policy, as described in Custom Security Intelligence Lists, on page 1438 and Configure Security Intelligence, on page 1888.

  For IP addresses, you can optionally use network objects rather than lists or feeds for this purpose; for information, see Network, on page 1401. (For URLs, using lists and feeds is strongly recommended over other methods.)

- Custom Do Not Block lists or feeds

Override Security Intelligence blocking for specific sites or addresses. See Override Security Intelligence Blocking, on page 1895.

- Global Block lists (one each for Network, URL and DNS)

  While reviewing events, you can immediately add an event's IP address, URL, or domain to the applicable Global Block List so that Security Intelligence will handle future traffic from that source. See Global and Domain Security Intelligence Lists, on page 1431.

- Global Do Not Block lists (one each for Network, URL and DNS)

  While reviewing events, you can immediately add an event's IP address, URL, or domain to the applicable Global Do Not Block List if you do not want Security Intelligence to block future traffic from that source. See Global and Domain Security Intelligence Lists, on page 1431.

Configure Security Intelligence

Each access control policy has Security Intelligence options. You can add network objects, URL objects and lists, and Security Intelligence feeds and lists to a Block list or Do Not Block list, and constrain any of these by security zone. You can also associate a DNS policy with your access control policy, and add domain names to a Block or Do Not Block list.

The number of objects in the Do Not Block lists plus the number in the Block lists cannot exceed 125 network objects, or 32767 URL objects and lists.

Before you begin

- Tip: For guidance on minimum configuration recommendations, see also Configuration Example: Security Intelligence Blocking, on page 1893.
- To ensure that all options are available to select, add at least one managed device to your management center.
- In passive deployments, or if you want to set Security Intelligence filtering to monitor-only, enable logging; see *Logging Connections with Security Intelligence* in the Cisco Secure Firewall Management Center Administration Guide .
- Configure a DNS policy to take Security Intelligence action for domains. For more information, see DNS Policies, on page 1897.

Procedure

Step 1 In the access control policy editor, click **Security Intelligence**.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 You have the following options:

- Click **Networks** to add network objects (IP addresses).

  Note

Network objects used in a Security Intelligence policy require a IPS license.

- Click **URLs** to add URL objects.

Step 3 Find the **Available Objects** you want to add to the Block or Do Not Block list. You have the following options:

- Search the available objects by typing in the **Search by name or value** field. Clear the search string by clicking **Reload** ( ) or **Clear** ( ).
- If no existing list or feed meets your needs, click **Add** ( ), select **New Network List** or **New URL List**, and proceed as described in Creating Security Intelligence Feeds, on page 1437 or Uploading New Security Intelligence Lists to the Secure Firewall Management Center, on page 1439.
- If no existing object meets your needs, click **Add** ( ), select **New Network Object** or **New URL Object**, and proceed as described in Creating Network Objects, on page 1403.

Security Intelligence ignores IP address blocks using a `/0` netmask.

Step 4 Choose one or more **Available Objects** to add.

Step 5 (Optional) Choose an **Available Zone** to constrain the selected objects by zone.

You cannot constrain system-provided Security Intelligence lists by zone.

Note
The **Any** zone for an SI list applies only to interfaces that are part of a security zone. However, an exception is that if a device does not have any interfaces associated with a security zone, then the **Any** zone will match any interface.

For example, if you have five interfaces on a device and none of them are associated with a security zone, any SI list that is assigned to the **Any** zone will be inspected against traffic on ALL interfaces on the device. If you add one interface to a security zone on that device, it effectively would remove SI inspection on the other four interfaces, where the zone is set to **Any** for an SI list. If you add the other four interfaces to a security zone, they will be evaluated by the SI list attached to the **Any** zone.

Step 6 Click **Add to Do Not Block list** or **Add to Block list**, or click and drag the selected objects to either list.

To remove an object from a Block or Do Not Block list, click **Delete** ( ) To remove multiple objects, choose the objects and right-click to **Delete Selected**.

Step 7 (Optional) Set objects on the Block list to monitor-only by right-clicking the object under **Block List**, then choosing **Monitor-only (do not block)**.

You cannot set system-provided global Security Intelligence lists to monitor only.

Step 8 Choose a DNS policy from the **DNS Policy** drop-down list.

Step 9 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Related Topics

Security Intelligence, on page 1429
Snort Restart Scenarios, on page 252

Security Intelligence Options

Use the Security Intelligence tab in the access control policy editor to configure network (IP address) and URL Security Intelligence, and to associate the access control policy with a DNS policy in which you have configured Security Intelligence for domains.

Available Objects

Available objects include:

- Security Intelligence categories populated by the system-provided feed.

  For details, see Security Intelligence Categories, on page 1891.

- System-provided Global Block and Do Not Block lists.

  For descriptions, see Security Intelligence Sources, on page 1887.

- Security Intelligence lists and feeds that you create under Object > Object Management > Security Intelligence.

  For descriptions, see Security Intelligence Sources, on page 1887.

- Network and URL objects and groups that are configured on the respective pages under Object > Object Management. These are different from the Security Intelligence objects in the previous bullet.

  For details about network objects, see Network, on page 1401. (For URLs, use Security Intelligence lists or feeds rather than objects or groups.)

Available Zones

Except for the system-provided Global lists, you can constrain Security Intelligence filtering by zone.

For example, to improve performance, you might want to target enforcement to specific zones rather than all interfaces. As a more specific example, you can block spam only for a security zone that handles email traffic.

To enforce Security Intelligence filtering for an object on multiple zones, you must add the object to the Block or Do Not Block list separately for each zone.

DNS Policy

In order to match DNS traffic using Security Intelligence, you must select a DNS policy for your Security Intelligence configuration.

Using Block or Do Not Block lists, or monitoring traffic based on a DNS list or feed, also requires that you:

- Configure DNS Security Intelligence lists and feeds. See Security Intelligence, on page 1429.
- Create a DNS policy. See Creating Basic DNS Policies, on page 1900 for more information.
- Configure DNS rules that reference your DNS lists or feeds. See Creating and Editing DNS Rules, on page 1902 for more information.

- Because you deploy the DNS policy as part of your access control policy, you must associate both policies. See DNS Policy Deploy, on page 1909 for more information.

Do Not Block List

See Override Security Intelligence Blocking, on page 1895.

To select all objects in the list, right-click an object.

Block List

See Configuration Example: Security Intelligence Blocking, on page 1893 and other topics in this chapter.

For explanations of the visual indicators in the Block list, see Block List Icons, on page 1893.

To select all objects in the list, right-click an object.

Logging

Security Intelligence logging, enabled by default, logs all blocked and monitored connections handled by an access control policy's assigned devices. However, the system does not log Do Not Block list matches; logging of connections on the Do Not Block list depends on their eventual disposition. Logging must be enabled for connections on the Block list before you can set objects on that list to monitor-only.

To enable, disable, or view logging settings, right-click an object in the Block list.

Related Topics

Global and Domain Security Intelligence Lists, on page 1431
Security Intelligence Lists and Multitenancy, on page 1431

Security Intelligence Categories

Security Intelligence categories are determined by the system-provided feeds described in Security Intelligence, on page 1429.

These categories are used in the following locations:

- The Networks sub-tab on the Security Intelligence tab of an access control policy
- The URLs sub-tab beside the Networks tab on the Security Intelligence tab of an access control policy
- In a DNS policy on the DNS tab in the DNS rule configuration page
- In events generated when traffic matches Block or Monitor configurations in the above locations

Note If your organization is using Secure Firewall Threat Intelligence Director: When viewing events, you may see categories that indicate that the action was taken by TID, such as `TID URL Block`.

Categories are updated by Talos from the cloud, and this list may change independently of Firepower releases.

Table 101: Cisco Talos Intelligence Group (Talos) Feed Categories

Security Intelligence Category	Description
Attackers	Active scanners and hosts known for outbound malicious activity
Banking_fraud	Sites that engage in fraudulent activities that relate to electronic banking
Bogon	Bogon networks and unallocated IP addresses
Bots	Sites that host binary malware droppers
CnC	Sites that host command-and-control servers for botnets
Cryptomining	Hosts providing remote access to pools and wallets for the purpose of mining cryptocurrency
Dga	Malware algorithms used to generate a large number of domain names acting as rendezvous points with their command-and-control servers
Exploitkit	Software kits designed to identify software vulnerabilities in clients
High_risk	Domains and hostnames that match against the OpenDNS predictive security algorithms from security graph
Ioc	Hosts that have been observed to engage in Indicators of Compromise (IOC)
Link_sharing	Websites that share copyrighted files without permission
Malicious	Sites exhibiting malicious behavior that do not necessarily fit into another, more granular, threat category
Malware	Sites that host malware binaries or exploit kits
Newly_seen	Domains that have recently been registered, or not yet seen via telemetry. **Attention** Currently, this category does not have any active feed and is reserved for future use.
Open_proxy	Open proxies that allow anonymous web browsing
Open_relay	Open mail relays that are known to be used for spam
Phishing	Sites that host phishing pages
Response	IP addresses and URLs that are actively participating in malicious or suspicious activity
Spam	Mail hosts that are known for sending spam
Spyware	Sites that are known to contain, serve, or support spyware and adware activities
Suspicious	Files that appear to be suspicious and have characteristics that resemble known malware

Security Intelligence Category	Description
Tor_exit_node	Hosts known to offer exit node services for the Tor Anonymizer network

Block List Icons

The following visual indicators may appear in the Block list on the Security Intelligence tab in an access control policy:

Icon or Visual Indicator	Description
Block ( )	The object is set to block.
Monitor ( )	The object is set to monitor-only. See Security Intelligence Monitoring, on page 1894
An object is displayed in strikethrough text	The same object is also on the Do Not Block list, which overrides the block.

Configuration Example: Security Intelligence Blocking

Configure your access control policy to block all threats detectable by the system's regularly updated Security Intelligence feeds.

The number of objects in the Block lists plus the number in the Do Not Block lists cannot exceed 125 network objects, or 32767 URL objects and lists.

Before you begin

- To ensure that all options are available to select, add at least one managed device to your management center.
- Configure a DNS policy to block all Security Intelligence threat categories for domains. For more information, see DNS Policies, on page 1897.
- If you have, or will have, custom lists of entities to block, create a Security Intelligence object of each type (URLs, DNS, Networks.) See Security Intelligence, on page 1429.

Procedure

Step 1 Click **Policies** > **Access Control heading** > **Access Control**.

Step 2 Create a new access control policy or edit an existing policy.

Step 3 In the access control policy editor, click **Security Intelligence**.

If the controls are dimmed, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 4 Click **Networks** to add blocking criteria for IP addresses.

a) Scroll down in the Networks list and select all of the threat categories listed below the Global lists.
b) If applicable, select the security zones for which you want to block these threats.
c) Click **Add to Block List**.
d) If you have created custom lists or feeds with addresses to block, add those to the Block List using the same steps as above.

Step 5 Click **URLs** to add blocking criteria for URLs, and repeat the steps you followed for Networks.

Step 6 Choose a DNS policy from the **DNS Policy** drop-down list; see DNS Policy Overview, on page 1897.

Step 7 Click **Save**.

What to do next

- Enable logging for these connections; see *Logging Connections with Security Intelligence* in the Cisco Secure Firewall Management Center Administration Guide .
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.
- For additional protection, configure URL filtering to block malicious URLs. See URL Filtering, on page 1857.

Security Intelligence Monitoring

Monitoring logs connection events for traffic that would have been blocked by Security Intelligence, but does not block the traffic. Monitoring is especially useful for:

- Testing feeds before you implement them.

  Consider a scenario where you want to test a third-party feed before you implement blocking using that feed. When you set the feed to monitor-only, the system allows connections that would have been blocked to be further analyzed by the system, but also logs a record of each of those connections for your evaluation.

- Passive deployments, to optimize performance.

  Managed devices that are deployed passively cannot affect traffic flow; there is no advantage to configuring the system to block traffic. Additionally, because blocked connections are not actually blocked in passive deployments, the system may report multiple beginning-of-connection events for each blocked connection.

Note If configured, Secure Firewall Threat Intelligence Director may impact the action taken (Monitor or Block.) For more information, see Threat Intelligence Director-Firewall Management Center Action Prioritization, on page 2757.

To Configure Security Intelligence Monitoring:

After you configure Security Intelligence blocking following the instructions in Configuration Example: Security Intelligence Blocking, on page 1893, right-click each applicable object in the Block list and choose **Monitor-only**. You cannot set system-provided Security Intelligence lists to monitor only.

Override Security Intelligence Blocking

Optionally, you can use Do Not Block lists to exempt specific domains, URLs, or IP addresses from being blocked by Security Intelligence lists or feeds.

For example, you can:

- Override the occasional false-positive block in a reputable Security Intelligence feed
- Inspect specific traffic in depth instead of blocking it early based on reputation
- Exempt otherwise-restricted transactions based on zone from Security Intelligence blocking

  For example, you can add an improperly classified URL to a Do Not Block list, but then restrict the Do Not Block list object using a security zone used by those in your organization who need to access those URLs. That way, only those with a business need can access the URLs on the Do Not Block list.

Note Entries on a Do Not Block list are simply exemptions from the block list. Any connection that passes the Security Intelligence policy is subject to the access control rules. Thus, an entry in the Do Not Block list can subsequently be blocked by an access control rule or intrusion policy. Your Do Not Block entries should always be exemptions from your block lists.

Procedure

Step 1 Option 1: Add an IP address, URL, or domain from an event to the Global Do Not Block List. See Global and Domain Security Intelligence Lists, on page 1431.

Step 2 Option 2: Use a custom Security Intelligence list or feed.

a) Create the custom Security Intelligence list or feed. See Custom Security Intelligence Lists, on page 1438 or Creating Security Intelligence Feeds, on page 1437.

b) For IP addresses (Networks) and URLs: Edit your access control policy, click the Security Intelligence tab, then click the custom list or feed in the Networks or URLs sub-tab, then click **Add to Do Not Block List**.

c) Save your changes.

d) For domains (DNS): See the "DNS Policy" section in the Security Intelligence Options, on page 1890 topic.

e) Deploy your changes.

Troubleshooting Security Intelligence

See the following topics for troubleshooting Security Intelligence.

Security Intelligence Categories Are Missing from the Available Options List

Symptoms: On the Security Intelligence tab of the access control policy, Security Intelligence categories (such as CnC or Exploitkit) are not displayed in the Networks tab under Available Options.

Cause:

- These categories do not appear until you have added at least one managed device to your management center. You must add a device in order to pull all TALOS feeds.
- The URL filtering feature uses a different set of categories than the Security Intelligence feature; the category that you expect to see may be a URL filtering category. To see URL filtering categories, look at the **URLs** tab in an access control rule.

History for Security Intelligence Block Listing

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
New Security Intelligence categories	All	Any	Talos has added the following new Security Intelligence categories: • banking_fraud • ioc • high_risk • link_sharing • malicious • newly_seen • spyware You should update your access control and DNS policies to address the new categories, and check periodically for future changes. New/modified pages: Security Intelligence tab, Networks and URLs sub-tabs; DNS rules in DNS policies Supported platforms: Firewall Management Center

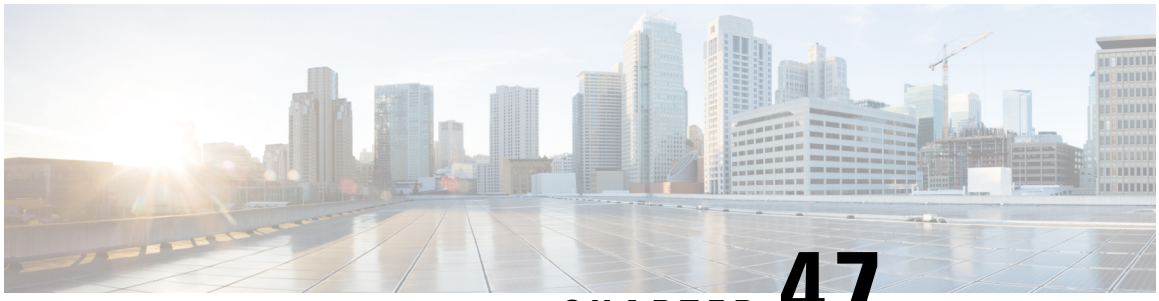

CHAPTER 47

DNS Policies

The following topics explain DNS policies, DNS rules, and how to deploy DNS policies to managed devices.

- DNS Policy Overview, on page 1897
- Cisco Umbrella DNS Policies, on page 1898
- DNS Policy Components, on page 1898
- License Requirements for DNS Policies, on page 1899
- Requirements and Prerequisites for DNS Policies, on page 1899
- Managing DNS and Umbrella DNS Policies, on page 1900
- DNS Rules, on page 1901
- How to Create DNS Rules, on page 1907
- DNS Policy Deploy, on page 1909
- Cisco Umbrella DNS Policies, on page 1909

DNS Policy Overview

DNS-based Security Intelligence allows you to block traffic based on the domain name requested by a client, using a Security Intelligence Block list. Cisco provides domain name intelligence you can use to filter your traffic; you can also configure custom lists and feeds of domain names tailored to your deployment.

Traffic on a DNS policy Block list is immediately blocked and therefore is not subject to any further inspection—not for intrusions, exploits, malware, and so on, but also not for network discovery. You can use a Security Intelligence Do Not Block list to override a Block list and force access control rule evaluation, and, recommended in passive deployments, you can use a "monitor-only" setting for Security Intelligence filtering. This allows the system to analyze connections that would have been blocked by a Block list, but also logs the match to the Block list and generates an end-of-connection Security Intelligence event.

Note DNS-based Security Intelligence may not work as intended for a domain name unless the DNS server deletes a domain cache entry due to expiration, or a client's DNS cache or the local DNS server's cache is cleared or expires.

You configure DNS-based Security Intelligence using a DNS policy and associated DNS rules. To deploy it to your devices, you must associate your DNS policy with an access control policy, then deploy your configuration to managed devices.

Cisco Umbrella DNS Policies

Cisco Umbrella DNS Connection in the management center helps to redirect DNS queries to Cisco Umbrella. This allows Cisco Umbrella to validate requests, allow or block them based on the domain names, and apply DNS-based security policy on the request. If you use Cisco Umbrella, you must configure the Cisco Umbrella Connection (**Integration > Other Integrations > Cloud Services > Cisco Umbrella Connection**) to redirect DNS queries to Cisco Umbrella.

The Umbrella Connector is part of the system's DNS inspection. If your existing DNS inspection policy map decides to block or drop a request based on your DNS inspection settings, the request is not forwarded to Cisco Umbrella. Thus, you have two lines of protection:

- Your local DNS inspection policy
- Your Cisco Umbrella cloud-based policy

When redirecting DNS lookup requests to Cisco Umbrella, the Umbrella Connector adds an EDNS (Extension mechanisms for DNS) record. An EDNS record includes the device identifier information, organization ID, and client IP address. Your cloud-based policy can use those criteria to control access in addition to the reputation of the FQDN. You can also elect to encrypt the DNS request using DNSCrypt to ensure the privacy of usernames and internal IP addresses.

To redirect DNS requests from the management center to Cisco Umbrella:

1. Configure the Cisco Umbrella connection settings.
2. Create and configure an Umbrella DNS policy.
3. Associate the Umbrella DNS policy with an access control policy.
4. Deploy the changes.

For detailed information about how to set up the Umbrella DNS Connector in the management center, see Configuring the Umbrella DNS Connector for Cisco Secure Firewall Management Center.

DNS Policy Components

A DNS policy allows you to block connections based on domain name, using a Block list, or exempt such connections from this type of blocking using a Do Not Block list. The following list describes the configurations you can change after creating a DNS policy.

Name and Description

Each DNS policy must have a unique name. A description is optional.

Rules

Rules provide a granular method of handling network traffic based on the domain name. Rules in a DNS policy are numbered, starting at 1. The system matches traffic to DNS rules in top-down order by ascending rule number.

When you create a DNS policy, the system populates it with a default Global Do-Not-Block List for DNS rule and a default Global Block List for DNS rule. Both rules are fixed to the first position in their respective categories. You cannot modify these rules, but you can disable them.

Note If multitenancy is enabled for your Firewall Management Center, the system is organized into a hierarchy of domains, including ancestor and descendant domains. These domains are distinct and separate from the domain names used in DNS management.

A descendant list contains the domains on the Block or Do Not Block lists of system subdomain users. From an ancestor domain, you cannot view the contents of descendant lists. If you do not want subdomain users to add domains to a Block or Do Not Block list:

- disable the descendant list rules, and
- enforce Security Intelligence using the access control policy inheritance settings

The system evaluates rules in the following order:

- Global Do-Not-Block List for DNS rule (if enabled)
- Descendant DNS Do-Not-Block Lists rule (if enabled)
- Rules with a Do Not Block action
- Global Block List for DNS rule (if enabled)
- Descendant DNS Block Lists rule (if enabled)
- Rules with an action other than Do Not Block

Usually, the system handles DN-based network traffic according to the *first* DNS rule where *all* the rule's conditions match the traffic. If no DNS rules match the traffic, the system continues evaluating the traffic based on the associated access control policy's rules. DNS rule conditions can be simple or complex.

License Requirements for DNS Policies

Threat Defense License

IPS

Requirements and Prerequisites for DNS Policies

Model Support

Any

Supported Domains

Any

User Roles

- Admin

- Access Admin
- Network Admin

Important You must apply the Network Discovery policy on the device for a successful DNS validation on the traffic.

Managing DNS and Umbrella DNS Policies

Use the DNS Policy page (**Policies** > **Access Control heading** > **DNS**) to manage custom DNS and Umbrella DNS policies.

In addition to custom policies that you create, the system provides the Default DNS Policy and Default Umbrella DNS Policy. The default DNS policy uses the default Block list and Do Not Block list. You can edit and use this system-provided custom policies.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **DNS**.

Step 2 Manage your DNS policy:

- Compare—To compare DNS policies, click **Compare Policies** and proceed as described in Compare Policies, on page 277.
- Copy—To copy a DNS policy, click **Copy** ( ) and proceed as described in Editing DNS Policies, on page 1901.
- Create—To create a new Umbrella DNS policy, click **New Policy** > **Umbrella DNS Policy** and proceed as described in Create an Umbrella DNS Policy, on page 1913.
- Delete—To delete a DNS or Umbrella DNS policy, click **Delete** ( ), then confirm you want to delete the policy.
- Edit—To modify an existing DNS policy, click **Edit** ( )) and proceed as described in Editing DNS Policies, on page 1901. To modify an existing Umbrella DNS policy, click **Edit** ( )) and proceed as described in Edit Umbrella DNS Policies and Rules, on page 1913.

Creating Basic DNS Policies

When you create a new DNS policy, it contains default settings. You must then edit it to customize the behavior.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **DNS**.

Step 2 Click **Add DNS Policy** > **DNS Policy**.

Step 3 Give the policy a unique **Name** and, optionally, a **Description**.

Step 4 Click **Save**.

What to do next

Configure the policy. See Editing DNS Policies, on page 1901.

Editing DNS Policies

Only one person should edit a DNS policy at a time, using a single browser window. If multiple users attempt to save the same policy, only the first set of saved changes are retained.

To protect the privacy of your session, after thirty minutes of inactivity on the policy editor, a warning appears. After sixty minutes, the system discards your changes.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **DNS**.

Step 2 Click **Edit** (✎) next to the DNS policy you want to edit.

If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Edit your DNS policy:

- Name and Description—To change the name or description, click the field and type the new information.
- Rules—To add, categorize, enable, disable, or otherwise manage DNS rules, click **Rules** and proceed as described in Creating and Editing DNS Rules, on page 1902.

Step 4 Click **Save**.

What to do next

- Optionally, further configure the new policy as described in *Logging Connections with Security Intelligence* in the Cisco Secure Firewall Management Center Administration Guide.
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

DNS Rules

DNS rules handle traffic based on the domain name requested by a host. As part of Security Intelligence, this evaluation happens after any traffic decryption, and before access control evaluation.

The system matches traffic to DNS rules in the order you specify. In most cases, the system handles network traffic according to the *first* DNS rule where *all* the rule's conditions match the traffic.

In addition to its unique name, each DNS rule has the following basic components:

State

By default, rules are enabled. If you disable a rule, the system does not use it to evaluate network traffic, and stops generating warnings and errors for that rule.

Position

Rules in a DNS policy are numbered, starting at 1. The system matches traffic to rules in top-down order by ascending rule number. With the exception of Monitor rules, the first rule that traffic matches is the rule that handles that traffic.

Conditions

Conditions specify the specific traffic the rule handles. A DNS rule must contain a DNS feed or list condition, and can also match traffic by security zone, network, or VLAN.

Action

A rule's action determines how the system handles matching traffic:

- Traffic with a **Do Not Block** action is allowed, subject to further access control inspection.
- Monitored traffic is subject to further evaluation by remaining rules on the DNS Block list. If the traffic does not match a DNS Block list rule, it is inspected with access control rules. The system logs a Security Intelligence event for the traffic.
- Traffic on a Block list is dropped without further inspection. You can also return a Domain Not Found response, or redirect the DNS query to a sinkhole server.

Related Topics

About Security Intelligence, on page 1885

Creating and Editing DNS Rules

In a DNS policy, you can add up to a total of 32767 DNS lists to the Block list and Do Not Block list rules; that is, the number of lists in the DNS policy cannot exceed 32767.

Procedure

Step 1 In the DNS policy editor, you have the following options:

- To add a new rule, click **Add DNS Rule**.
- To edit an existing rule, click **Edit** ( ).

Step 2 Enter a **Name**.

Step 3 Configure the rule components, or accept the defaults:

- Action—Choose a rule **Action**; see DNS Rule Actions, on page 1904.
- Conditions—Configure the rule's conditions; see DNS Rule Conditions, on page 1905.

- Enabled—Specify whether the rule is **Enabled**.

Step 4 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

DNS Rule Management

The **Rules** tab of the DNS policy editor allows you to add, edit, move, enable, disable, delete, and otherwise manage DNS rules within your policy.

For each rule, the policy editor displays its name, a summary of its conditions, and the rule action. Other icons represent **Warning** ( ), **Error** ( ), and other important **Information** ( ). Disabled rules are dimmed and marked `(disabled)` beneath the rule name.

Enabling and Disabling DNS Rules

When you create a DNS rule, it is enabled by default. If you disable a rule, the system does not use it to evaluate network traffic and stops generating warnings and errors for that rule. When viewing the list of rules in a DNS policy, disabled rules are dimmed, although you can still modify them. Note that you can also enable or disable a DNS rule using the DNS rule editor.

Procedure

Step 1 In the DNS policy editor, right-click the rule and choose a rule state.

Step 2 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

DNS Rule Order Evaluation

Rules in a DNS policy are numbered, starting at 1. The system matches traffic to DNS rules in top-down order by ascending rule number. In most cases, the system handles network traffic according to the *first* DNS rule where *all* the rule's conditions match the traffic:

- For Monitor rules, the system logs the traffic, then continues evaluating traffic against lower-priority DNS Block list rules.
- For non-Monitor rules, the system does **not** continue to evaluate traffic against additional, lower-priority DNS rules after that traffic matches a rule.

Note the following regarding rule order:

- The Global Do-Not-Block List for DNS is always first, and takes precedence over all other rules.
- The Do-Not-Block List section precedes the Block List section; Do-Not-Block List rules always take precedence over other rules.
- The Global Block List for DNS is always first in the Block List section, and takes precedence over all other Monitor and Block list rules.
- The Block List section contains Monitor and Block list rules.
- When you first create a DNS rule, the system positions it last in the Do-Not-Block List section if you assign a **Do Not Block** action, or last in the Block List section if you assign any other action.

You can drag and drop rules to reorder them.

DNS Rule Actions

Every DNS rule has an *action* that determines the following for matching traffic:

- handling—foremost, the rule action governs whether the system will block, not block, or monitor traffic that matches the rule's conditions, based on a Block or Do Not Block list
- logging—the rule action determines when and how you can log details about matching traffic

If configured, TID also impacts action prioritization. For more information, see Threat Intelligence Director-Firewall Management Center Action Prioritization, on page 2757.

Do Not Block Action

The **Do Not Block** action allows traffic to pass to the next phase of inspection, which is access control rules.

The system does not log Do Not Block list matches. Logging of these connections depends on their eventual disposition.

Monitor Action

The **Monitor** action is designed to force connection logging; matching traffic is neither immediately allowed nor blocked. Rather, traffic is matched against additional rules to determine whether to permit or deny it. The first non-Monitor DNS rule matched determines whether the system blocks the traffic. If there are no additional matching rules, the traffic is subject to access control evaluation.

For connections monitored by a DNS policy, the system logs end-of-connection Security Intelligence and connection events to the Firewall Management Center database.

Block Actions

These actions block traffic without further inspection of any kind:

- The **Drop** action drops the traffic.
- The **Domain Not Found** action returns a non-existent internet domain response to the DNS query, which prevents the client from resolving the DNS request.
- The **Sinkhole** action returns a sinkhole object's IPv4 or IPv6 address in response to the DNS query (A and AAAA records only). The sinkhole server can log, or log and block, follow-on connections to the IP address. If you configure a **Sinkhole** action, you must also configure a sinkhole object.

For a connection blocked based on the **Drop** or **Domain Not Found** actions, the system logs beginning-of-connection Security Intelligence and connection events. Because blocked traffic is immediately denied without further inspection, there is no unique end of connection to log.

For a connection blocked based on the **Sinkhole** action, logging depends on the sinkhole object configuration. If you configure your sinkhole object to only log sinkhole connections, the system logs end-of-connection connection events for the follow-on connection. If you configure your sinkhole object to log and block sinkhole connections, the system logs beginning-of-connection connection events for the follow-on connection, then blocks that connection.

DNS Rule Conditions

A DNS rule's conditions identify the type of traffic that rule handles. Conditions can be simple or complex. You must define a DNS feed or list condition within a DNS rule. You can also optionally control traffic by security zone, network, or VLAN.

When adding conditions to a DNS rule:

- If you do not configure a particular condition for a rule, the system does not match traffic based on that criterion.
- You can configure multiple conditions per rule. Traffic must match **all** the conditions in the rule for the rule to apply to traffic. For example, a rule with a DNS feed or list condition and network condition but no VLAN tag condition evaluates traffic based on the domain name and source or destination, regardless of any VLAN tagging in the session.
- For each condition in a rule, you can add up to 50 criteria. Traffic that matches **any** of a condition's criteria satisfies the condition. For example, you can use a single rule to block traffic based on up to 50 DNS lists and feeds.

The following topics explain DNS rule conditions in more detail.

Security Zone Rule Conditions

Security zones segment your network to help you manage, classify, and decrypt traffic flow by grouping interfaces across multiple devices.

Security zones control or decrypt traffic by its source and destination security zones. If you add both source and destination zones to a zone condition, matching traffic must originate from an interface in one of the source zones and leave through an interface in one of the destination zones.

Just as all interfaces in a zone must be of the same type (all inline, passive, switched, or routed), all zones used in a zone condition must be of the same type. Because devices deployed passively do not transmit traffic, you cannot use a zone with passive interfaces as a destination zone.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Tip Constraining rules by zone is one of the best ways to improve system performance. If a rule does not apply to traffic through any of device's interfaces, that rule does not affect that device's performance.

Security Zone Conditions and Multitenancy

In a multidomain deployment, a zone created in an ancestor domain can contain interfaces that reside on devices in different domains. When you configure a zone condition in an descendant domain, your configurations apply to only the interfaces you can see.

Network Rule Conditions

Networks control or decrypt traffic by its source and destination IP address, using inner headers. Tunnel rules, which use outer headers, have tunnel endpoint conditions instead of network conditions.

You can use predefined objects to build network conditions, or manually specify individual IP addresses or address blocks.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Note You *cannot* use FDQN network objects in identity rules.

VLAN Tags Rule Conditions

Note VLAN tags in access rules only apply to inline sets. Access rules with VLAN tags do not match traffic on firewall interfaces.

VLAN rule conditions control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic. The system uses the innermost VLAN tag to filter VLAN traffic, with the exception of the prefilter policy, which uses the outermost VLAN tag in its rules.

Note the following Q-in-Q support:

- Firewall Threat Defense on Firepower 4100/9300—Does not support Q-in-Q (supports only one VLAN tag).
- Firewall Threat Defense on all other models:
  - Inline sets and passive interfaces—Supports Q-in-Q, up to 2 VLAN tags.
  - Firewall interfaces—Does not support Q-in-Q (supports only one VLAN tag).

You can use predefined objects to build VLAN conditions, or manually enter any VLAN tag from 1 to 4094. Use a hyphen to specify a range of VLAN tags.

In a cluster, if you encounter problems with VLAN matching, edit the access control policy advanced options, Transport/Network Preprocessor Settings, and select the **Ignore the VLAN header when tracking connections** option.

DNS Policy Rule Conditions

DNS conditions in DNS rules allow you to control traffic if a DNS list, feed, or category contains the domain name requested by the client. You must define a DNS condition in a DNS rule.

Regardless of whether you add a global or custom Block or Do Not Block list to a DNS condition, the system applies the configured rule action to the traffic. For example, if you add the Global Do Not Block List to a rule, and configure a **Drop** action, the system blocks all traffic that should have been allowed to pass to the next phase of inspection.

How to Create DNS Rules

The following topics discuss how to create DNS rules.

Controlling Traffic Based on DNS and Security Zone

Zone conditions in DNS rules allow you to control traffic by its source security zone. A *security zone* is a grouping of one or more interfaces, which may be located across multiple devices.

Procedure

Step 1 In the DNS rule editor, click **Zones**.

Step 2 Find and select the zones you want to add from the **Available Zones**. To search for zones to add, click the **Search by name** prompt above the **Available Zones** list, then type a zone name. The list updates as you type to display matching zones.

Step 3 Click to select a zone, or right-click and then select **Select All**.

Step 4 Click **Add to Source**, or drag and drop.

Step 5 Click the **DNS** tab and add the lists or feeds that include the DNS names you are controlling. For more information, see Controlling Traffic Based on DNS List or Feed, on page 1909.

Step 6 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Controlling Traffic Based on DNS and Network

Network conditions in DNS rules allow you to control traffic by its source IP address. You can explicitly specify the source IP addresses for the traffic you want to control.

Procedure

Step 1 In the DNS rule editor, click **Networks**.

Step 2 Find and select the networks you want to add from the **Available Networks**, as follows:

- To add a network object on the fly, which you can then add to the condition, click **Add** (+) above the **Available Networks** list and proceed as described in Creating Network Objects, on page 1403.

- To search for network objects to add, click the **Search by name or value** prompt above the **Available Networks** list, then type an object name or the value of one of the object's components. The list updates as you type to display matching objects.

Step 3 Click **Add to Source**, or drag and drop.

Step 4 Add any source IP addresses or address blocks that you want to specify manually. Click the **Enter an IP address** prompt below the **Source Networks** list; then type an IP address or address block and click **Add**.

Step 5 Click the **DNS** tab and add the lists or feeds that include the DNS names you are controlling. For more information, see Controlling Traffic Based on DNS List or Feed, on page 1909.

Step 6 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Controlling Traffic Based on DNS and VLAN

VLAN conditions in DNS rules allow you to control VLAN-tagged traffic. The system uses the innermost VLAN tag to identify a packet by VLAN.

When you build a VLAN-based DNS rule condition, you can manually specify VLAN tags. Alternately, you can configure VLAN conditions with VLAN tag *objects*, which are reusable and associate a name with one or more VLAN tags.

Procedure

Step 1 In the DNS rule editor, select **VLAN Tags**.

Step 2 Find and select the VLANs you want to add from the **Available VLAN Tags**, as follows:

- To add a VLAN tag object on the fly, which you can then add to the condition, click **Add** (+) above the Available VLAN Tags list and proceed as described in Creating VLAN Tag Objects, on page 1461.
- To search for VLAN tag objects and groups to add, click the **Search by name or value** prompt above the **Available VLAN Tags** list, then type either the name of the object, or the value of a VLAN tag in the object. The list updates as you type to display matching objects.

Step 3 Click **Add to Rule**, or drag and drop.

Step 4 Add any VLAN tags that you want to specify manually. Click the **Enter a VLAN Tag** prompt below the **Selected VLAN Tags** list; then type a VLAN tag or range and click **Add**. You can specify any VLAN tag from 1 to 4094; use a hyphen to specify a range of VLAN tags.

Step 5 Click the **DNS** tab and add the lists or feeds that include the DNS names you are controlling. For more information, see Controlling Traffic Based on DNS List or Feed, on page 1909.

Step 6 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Controlling Traffic Based on DNS List or Feed

Procedure

Step 1 In the DNS rule editor, click **DNS**.

Step 2 Find and select the DNS lists and feeds you want to add from the **DNS Lists and Feeds**, as follows:

- To add a DNS list or feed on the fly, which you can then add to the condition, click **Add** (+) above the **DNS Lists and Feeds** list and proceed as described in Creating Security Intelligence Feeds, on page 1437 .
- To search for DNS lists, feeds, or categories to add, click the **Search by name or value** prompt above the **DNS Lists and Feeds** list, then type an object name or the value of one of the object's components. The list updates as you type to display matching objects.
- For descriptions of the system-provided threat categories, see Security Intelligence Categories, on page 1891.

Step 3 Click **Add to Rule**, or drag and drop.

Step 4 Save or continue editing the rule.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

DNS Policy Deploy

After you finish updating your DNS policy configuration, you must deploy it as part of access control configuration.

- Associate your DNS policy with an access control policy, as described in Configure Security Intelligence, on page 1888.
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Cisco Umbrella DNS Policies

Cisco Umbrella DNS Connection in the management center helps to redirect DNS queries to Cisco Umbrella. This allows Cisco Umbrella to validate requests, allow or block them based on the domain names, and apply DNS-based security policy on the request. If you use Cisco Umbrella, you must configure the Cisco Umbrella

Connection (**Integration > Other Integrations > Cloud Services > Cisco Umbrella Connection**) to redirect DNS queries to Cisco Umbrella.

The Umbrella Connector is part of the system's DNS inspection. If your existing DNS inspection policy map decides to block or drop a request based on your DNS inspection settings, the request is not forwarded to Cisco Umbrella. Thus, you have two lines of protection:

- Your local DNS inspection policy
- Your Cisco Umbrella cloud-based policy

When redirecting DNS lookup requests to Cisco Umbrella, the Umbrella Connector adds an EDNS (Extension mechanisms for DNS) record. An EDNS record includes the device identifier information, organization ID, and client IP address. Your cloud-based policy can use those criteria to control access in addition to the reputation of the FQDN. You can also elect to encrypt the DNS request using DNSCrypt to ensure the privacy of usernames and internal IP addresses.

To redirect DNS requests from the management center to Cisco Umbrella:

1. Configure the Cisco Umbrella connection settings.
2. Create and configure an Umbrella DNS policy.
3. Associate the Umbrella DNS policy with an access control policy.
4. Deploy the changes.

For detailed information about how to set up the Umbrella DNS Connector in the management center, see Configuring the Umbrella DNS Connector for Cisco Secure Firewall Management Center.

How to Redirect DNS Requests to Cisco Umbrella

This section provides instructions to redirect DNS requests from the device to Cisco Umbrella using the Firewall Management Center.

Step	Do This	More Info
1	Ensure that you meet the prerequisites.	Prerequisites for Configuring the Umbrella DNS Connector, on page 1911
2	Configure the Cisco Umbrella connection settings.	Configure Cisco Umbrella Connection Settings, on page 1912
3	Create an Umbrella DNS policy	Create an Umbrella DNS Policy, on page 1913
4	Configure the Umbrella DNS policy	Edit Umbrella DNS Policies and Rules, on page 1913
5	Associate the Umbrella DNS policy with an access control policy	Associate the Umbrella DNS Policy with an Access Control Policy, on page 1914

Prerequisites for Configuring the Umbrella DNS Connector

Table 102: Minimum Supported Platforms

Product	Version
Secure Firewall Threat Defense	6.6 and above
Secure Firewall Management Center	7.2 and above

- Establish an account with Cisco Umbrella at https://umbrella.cisco.com, and log into Umbrella at http://login.umbrella.com.
- Import the CA certificate from the Cisco Umbrella server to the Firewall Management Center. In Cisco Umbrella, choose **Deployments > Configuration > Root Certificate** and download the certificate.

  You must import the root certificate to establish the HTTPS connection with the Cisco Umbrella registration server. The certificate needs to be trusted for SSL server validation, which is a non-default option in the Firewall Management Center. Copy and paste the following certificate for the device in the Firewall Management Center(**Devices** > **Certificates**).

```
MIIE6jCCA9KgAwIBAgIQCjUI1VwpKwF9+K1lwA/35DANBgkqhkiG9w0BAQsFADBhMQswCQYDVQQG
EwJVUzEVMBMGA1UEChMMRGlnaUNlcnQgSW5jMRkwFwYDVQQLExB3d3cuZGlnaWNlcnQuY29tMSAw
HgYDVQQDExdEaWdpQ2VydCBHbG9iYWwgUm9vdCBDQTAeFw0yMDA5MjQwMDAwMDBaFw0zMDA5MjMy
MzU5NTlaME8xCzAJBgNVBAYTAlVTMRUwEwYDVQQKEwxEaWdpQ2VydCBJbmMxKTAnBgNVBAMTIERp
Z2lDZXJ0IFRMUyBSU0EgU0hBMjU2IDIwMjAgQ0ExMIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIB
CgKCAQEAwUuzZUdwvN1PWNvsnO3DZuUfMRNUrUpmRh8sCuxkB+Uu3Ny5CiDt3+PE0J6aqXodgojl
EVbbHp9YwlHnLDQNLtKS4VbL8Xlfs7uHyiUDe5pSQWYQYE9XE0nw6Ddng9/n00tnTCJRpt8OmRDt
V1F0JuJ9x8piLhMbfyOIJVNvwTRYAIuE//i+p1hJInuWraKImxW8oHzf6VGo1bDtN+I2tIJLYrVJ
muzHZ9bjPvXj1hJeRPG/cUJ9WIQDgLGBAfr5yjK7tI4nhyfFK3TUqNaX3sNk+crOU6JWvHgXjkkD
Ka77SU+kFbnO8lwZV21reacroicgE7XQPUDTITAHk+qZ9QIDAQABo4IBrjCCAaowHQYDVR0OBBYE
FLdrouqoqoSMeeq02g+YssWVdrn0MB8GA1UdIwQYMBaAFAPeUDVW0Uy7ZvCj4hsbw5eyPdFVMA4G
A1UdDwEB/wQEAwIBhjAdBgNVHSUEFjAUBggrBgEFBQcDAQYIKwYBBQUHAwIwEgYDVR0TAQH/BAgw
BgEB/wIBADB2BggrBgEFBQcBAQRqMGgwJAYIKwYBBQUHMAGGGGh0dHA6Ly9vY3NwLmRpZ2ljZXJ0
LmNvbTBABggrBgEFBQcwAoY0aHR0cDovL2NhY2VydHMuZGlnaWNlcnQuY29tL0RpZ2lDZXJ0R2xv
YmFsUm9vdENBLmNydDB7BgNVHR8EdDByMDegNaAzhjFodHRwOi8vY3JsMy5kaWdpY2VydC5jb20v
RGlnaUNlcnRHbG9iYWxSb290Q0EuY3JsMDegNaAzhjFodHRwOi8vY3JsNC5kaWdpY2VydC5jb20v
RGlnaUNlcnRHbG9iYWxSb290Q0EuY3JsMDAGA1UdIAQpMCcwBwYFZ4EMAQEwCAYGZ4EMAQIBMAgG
BmeBDAECAjAIBgZngQwBAgMwDQYJKoZIhvcNAQELBQADggEBAHert3onPa679n/gWlbJhKrKW3EX
3SJH/E6f7tDBpATho+vFScH90cnfjK+URSxGKqNjOSD5nkoklEHIqdninFQFBstcHL4AGw+oWv8Z
u2XHFq8hVt1hBcnpj5h232sb0HIMULkwKXq/YFkQZhM6LawVEWwtIwwCPgU7/uWhnOKK24fXSuhe
50gG66sSmvKvhMNbg0qZgYOrAKHKCjxMoiWJKiKnpPMzTFuMLhoClw+dj20tlQj7T9rxkTgl4Zxu
YRiHas6xuwAwapu3r9rxxZf+ingkquqTgLozZXq8oXfpf2kUCwA/d5KxTVtzhwoT0JzI8ks5T1KE
SaZMkE4f97Q=
```

  When you add the certificate in the management center, ensure that you check the **CA Only** check box.
- Install the certificate on the device.
- Obtain the following data from Umbrella:
  - Organization ID
  - Network Device Key
  - Network Device Secret
  - Legacy Network Device Token
- Ensure that the Firewall Management Center is connected to the internet.

- Ensure that the base license with the export-controlled feature option is enabled in the Firewall Management Center.
- Ensure that the DNS server is configured to resolve api.opendns.com.
- Ensure that the Firewall Management Center can resolve management.api.umbrella.com for policy configuration.
- Configure the Firewall Threat Defense route to api.opendns.com.

Configure Cisco Umbrella Connection Settings

The Cisco Umbrella Connection settings define the token that is needed to register the device with Cisco Umbrella.

Before you begin

Establish an account with Cisco Umbrella https://umbrella.cisco.com, and then log into Umbrella at https://dashboard.umbrella.com and obtain the required information to establish connection to Cisco Umbrella.

Procedure

Step 1 Choose **Integration** > **Other Integrations** > **Cloud Services** > **Cisco Umbrella Connection**.

Step 2 Obtain the following details and add them to the **General** settings:

- **Organization ID**—A unique number that identifies your organization on Cisco Umbrella. Every Umbrella organization is a separate instance of Umbrella and has its own dashboard. Organizations are identified by their name and their organization ID (Org ID).
- **Network Device Key**—The key to fetch umbrella policy from Cisco Umbrella.
- **Network Device Secret** —The secret to fetch umbrella policy from Cisco Umbrella.
- **Legacy Network Device Token** —An Umbrella Legacy Network Devices API token is issued through the Cisco Umbrella dashboard. Umbrella requires the API token to register a network device.

Step 3 Under **Advanced**, configure the following optional settings:

- **DNSCrypt Public Key**—DNSCrypt authenticates and encrypts the DNS queries between the endpoint and the DNS server. To enable DNScrypt, you can configure the DNScrypt public key for certificate verification. The key is a 32-byte hexadecimal value and is preconfigured to B735:1140:206F:225d:3E2B:d822:D7FD:691e:A1C3:3cc8:D666:8d0c:BE04:bfab:CA43:FB79, which is the public key of the Umbrella Anycast servers.
- **Management Key**—A key to fetch datacenter details from Umbrella cloud for VPN policy.
- **Management Secret**—A secret used to fetch datacenters from Umbrella cloud for VPN.

Step 4 Click **Test Connection**—Test if the Cisco Umbrella Cloud is reachable from the Firewall Management Center. When you provide the required organization ID and network device details, the umbrella connection is created.

Step 5 Click **Save**.

Create an Umbrella DNS Policy

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **DNS**.

Step 2 Click **Add DNS Policy** > **Umbrella DNS Policy**.

Step 3 Give the policy a unique **Name** and, optionally, a **Description**.

Step 4 Click **Save**.

What to do next

Configure the policy. See Edit Umbrella DNS Policies and Rules, on page 1913.

Edit Umbrella DNS Policies and Rules

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **DNS**.

Step 2 In the DNS Policy page, select the Umbrella DNS policy that you want to edit and click **Edit** ( ).

Refresh the Umbrella Protection Policy

If you want to get the latest Umbrella Protection Policy from Cisco Umbrella, click the **Refresh** icon next to **Umbrella Protection Policy Last Updated**.

To configure or modify Umbrella Connection settings for the Management Center, go to **Integration** > **Other Integrations** > **Cloud Services** > **Cisco Umbrella Connection**.

Step 3 In the Umbrella DNS policy editor, select the Umbrella DNS rule and click **Edit** ( ).

Step 4 Configure the rule components, or accept the defaults:

- **Umbrella Protection Policy**—Specify the name of the Cisco Umbrella policy to apply to the device.
- **Bypass Domain**—Specify the name of the local domains for which DNS requests should bypass Cisco Umbrella and instead go directly to the configured DNS servers.

  For example, you can have your internal DNS server resolve all names for the organization's domain name on the assumption that all internal connections are allowed.
- **DNSCrypt**— Enable DNSCrypt to encrypt connections between the device and Cisco Umbrella.

  Enabling DNScrypt starts the key-exchange thread with the Umbrella resolver. The key-exchange thread performs the handshake with the resolver every hour and updates the device with a new secret key. As

DNScrypt uses UDP/443, you must ensure that the class map used for DNS inspection includes that port. Note that the default inspection class already includes UDP/443 for DNS inspection.

- **Idle Timeout**—Configure the idle timeout after which a connection from a client to the Umbrella server will be removed if there is no response from the server.

Step 5 Click **Save**.

What to do next

Associate the Umbrella DNS policy with an access control policy. For more information, see Associate the Umbrella DNS Policy with an Access Control Policy, on page 1914.

Associate the Umbrella DNS Policy with an Access Control Policy

Before you deploy the Umbrella DNS policy on the device, you must associate it with an access control policy.

Procedure

Step 1 Choose **Policies** > **Access Control** and select the access policy to edit.

Step 2 Select **Security Intelligence**.

Step 3 From the **Umbrella DNS Policy** drop-down list, select the Umbrella DNS policy.

Step 4 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 259.

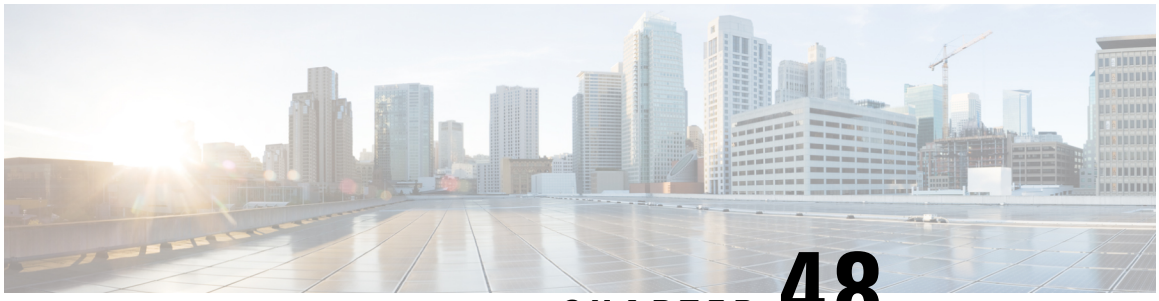

CHAPTER 48

Prefiltering and Prefilter Policies

- About Prefiltering, on page 1915
- Best Practices for Fastpath Prefiltering, on page 1920
- Best Practices for Encapsulated Traffic Handling , on page 1920
- Requirements and Prerequisites for Prefilter Policies, on page 1921
- Configure Prefiltering, on page 1922
- Tunnel Zones and Prefiltering, on page 1928
- Moving Prefilter Rules to an Access Control Policy, on page 1931
- Prefilter Policy Hit Counts, on page 1933
- Large Flow Offloads, on page 1933
- History for Prefiltering, on page 1936

About Prefiltering

Prefiltering is the first phase of access control, before the system performs more resource-intensive evaluation. Prefiltering is simple, fast, and early. Prefiltering uses limited outer-header criteria to quickly handle traffic. Compare this to subsequent evaluation, which uses inner headers and has more robust inspection capabilities.

Configure prefiltering to:

- Improve performance— The sooner you exclude traffic that does not require inspection, the better. You can fastpath or block certain types of plaintext, passthrough tunnels based on their outer encapsulation headers, without inspecting their encapsulated connections. You can also fastpath or block any other connections that benefit from early handling.
- Tailor deep inspection to encapsulated traffic—You can rezone certain types of tunnels so that you can later handle their encapsulated connections using the same inspection criteria. Rezoning is necessary because after prefiltering, access control uses inner headers.

About Prefilter Policies

Prefiltering is a policy-based feature. To assign it to a device, you assign it to the access control policy that is assigned to the device.

Policy Components: Rules and Default Action

In a prefilter policy, *tunnel rules*, *prefilter rules*, and a *default action* handle network traffic:

- Tunnel and prefilter rules—First, rules in a prefilter policy handle traffic in the order you specify. Tunnel rules match specific tunnels only and support rezoning. Prefilter rules have a wider range of constraints and do not support rezoning. For more information, see Tunnel vs Prefilter Rules, on page 1916.
- Default action (tunnels only)—If a tunnel does not match any rules, the default action handles it. The default action can block these tunnels, or continue access control on their individual encapsulated connections. You cannot rezone tunnels with the default action.

  There is no default action for nonencapsulated traffic. If a nonencapsulated connection does not match any prefilter rules, the system continues with access control.

Connection Logging

You can log connections fastpathed and blocked by the prefilter policy. See *Other Connections You Can Log* in the Cisco Secure Firewall Management Center Administration Guide for more information.

Connection events contain information on whether and how logged connections—including entire tunnels—were prefiltered. You can view this information in event views (workflows), dashboards, and reports, and use it as correlation criteria. Keep in mind that because fastpathed and blocked connections are not subject to deep inspection, associated connection events contain limited information.

Default Prefilter Policy

Every access control policy has an associated prefilter policy.

The system uses a default policy if you do not configure custom prefiltering. Initially, this system-provided policy passes all traffic to the next phase of access control. You can change the policy's default action and configure its logging options, but you cannot add rules to it or delete it.

Prefilter Policy Inheritance and Multitenancy

Access control uses a hierarchical implementation that complements multitenancy. Along with other advanced settings, you can lock a prefilter policy association, enforcing that association in all descendant access control policies. For more information, see Access Control Policy Inheritance, on page 1731.

Tunnel vs Prefilter Rules

Whether you configure a tunnel or prefilter rule depends on the specific type of traffic you want to match and the actions or further analysis you want to perform.

Characteristic	Tunnel Rules	Prefilter Rules
Primary function	Quickly fastpath, block, or rezone plaintext, passthrough tunnels.	Quickly fastpath or block any other connection that benefits from early handling.
Encapsulation and port/protocol criteria	Encapsulation conditions match only plaintext tunnels over selected protocols, listed in Encapsulation Rule Conditions, on page 1927.	Port conditions can use a wider range of port and protocol constraints than tunnel rules; see Port, Protocol, and ICMP Code Rule Conditions, on page 941.

Characteristic	Tunnel Rules	Prefilter Rules
Network criteria	Tunnel endpoint conditions constrain the endpoints of the tunnels you want to handle; see Network Rule Conditions, on page 939.	Network conditions constrain the source and destination hosts in each connection; see Network Rule Conditions, on page 939.
Direction	Bidirectional or unidirectional (configurable). Tunnel rules are bidirectional by default, so they can handle all traffic between tunnel endpoints.	Unidirectional only (nonconfigurable). Prefilter rules match source-to-destination traffic only. Return traffic for allowed connections is also permitted.
Rezone sessions for further analysis	Supported, using tunnel zones; see Tunnel Zones and Prefiltering, on page 1928.	Not supported.

Prefiltering vs Access Control

Prefilter and access control policies both allow you to block and trust traffic, though the prefiltering "trust" functionality is called "fastpathing" because it skips more inspection. The following table explains this and other differences between prefiltering and access control, to help you decide whether to configure custom prefiltering.

If you do not configure custom prefiltering, you can only approximate—not replicate—prefilter functionality with early-placed Block and Trust rules in the access control policy.

Characteristic	Prefiltering	Access Control	For more information, see...
Primary function	Quickly fastpath or block certain types of plaintext, passthrough tunnels (see Encapsulation Rule Conditions, on page 1927), or tailor subsequent inspection to their encapsulated traffic. Fastpath or block any other connections that benefit from early handling.	Inspect and control all network traffic, using simple or complex criteria, including contextual information and deep inspection results.	About Prefiltering, on page 1915
Implementation	Prefilter policy. The prefilter policy is invoked by the access control policy.	Access control policy. The access control policy is a main configuration. In addition to invoking subpolicies, access control policies have their own rules.	About Prefilter Policies, on page 1915 Associating Other Policies with Access Control, on page 1762

Characteristic	Prefiltering	Access Control	For more information, see...
Sequence within access control	First. The system matches traffic to prefilter criteria before all other access control configurations. Thus, fast-pathing in a prefilter policy skips many other policies and inspections.	Last. Other policies that might inspect a connection, including Security Intelligence and Decryption, are applied before access control rules are evaluated.	—
Rule actions	Fewer. You can stop further inspection (Fastpath and Block) or allow further analysis with the rest of access control (Analyze).	More. Access control rules have a larger variety of actions, including monitoring, deep inspection, block with reset, and interactive blocking.	Tunnel and Prefilter Rule Components, on page 1923 Access Control Rule Actions, on page 1778
Bypass capability	Fastpath rule action. Fastpathing traffic in the prefilter stage bypasses all further inspection and handling, including: • Security Intelligence • authentication requirements imposed by an identity policy • SSL decryption • access control rules • deep inspection of packet payloads • discovery • rate limiting	Trust rule action. Traffic trusted by access control rules is only exempt from intrusion inspection and discovery. If another policy that requires inspection is applied to a connection that is eventually trusted, the connection is inspected anyway. Trust really just means you are not applying intrusion or file policies to the connection.	Introduction to Access Control Rules, on page 1773
Rule criteria	Limited. Rules in the prefilter policy use simple network criteria: IP address, VLAN tag, port, and protocol. For tunnels, tunnel endpoint conditions specify the IP address of the routed interfaces of the network devices on either side of the tunnel.	Robust. Access control rules use network criteria, but also user, application, requested URL, and other contextual information available in packet payloads. Network conditions specify the IP address of source and destination hosts.	Tunnel vs Prefilter Rules, on page 1916 Prefilter Rule Conditions, on page 1925 Tunnel Rule Conditions, on page 1927

Characteristic	Prefiltering	Access Control	For more information, see...
IP headers used (tunnel handling)	Outermost. Using outer headers allows you to handle entire plaintext, passthrough tunnels. For nonencapsulated traffic, prefiltering still uses "outer" headers—which in this case are the only headers.	Innermost possible. For a nonencrypted tunnel, access control acts on its individual encapsulated connections, not the tunnel as a whole.	Passthrough Tunnels and Access Control, on page 1919
Rezone encapsulated connections for further analysis	Rezones tunneled traffic. Tunnel zones allow you to tailor subsequent inspection to prefiltered, encapsulated traffic.	Uses tunnel zones. Access control uses the tunnel zones you assign during prefiltering.	Tunnel Zones and Prefiltering, on page 1928
Connection logging	Fastpathed and blocked traffic only. Allowed connections may still be logged by other configurations.	Any connection.	*Other Connections You Can Log* in the Cisco Secure Firewall Management Center Administration Guide
Supported devices	Secure Firewall Threat Defense only.	All.	—

Passthrough Tunnels and Access Control

Plaintext (nonencrypted) tunnels can encapsulate multiple connections, often flowing between discontinuous networks. These tunnels are especially useful for routing custom protocols over IP networks, IPv6 traffic over IPv4 networks, and so on.

An outer *encapsulation header* specifies the source and destination IP addresses of the *tunnel endpoints*—the routed interfaces of the network devices on either side of the tunnel. Inner *payload headers* specify the source and destination IP addresses of the encapsulated connections' actual endpoints.

Often, network security devices handle plaintext tunnels as *passthrough* traffic. That is, the device is not one of the tunnel endpoints. Instead, it is deployed between the tunnel endpoints and monitors the traffic flowing between them.

Some network security devices enforce security policies using outer IP headers. Even for plaintext tunnels, these devices have no control over or insight into individual encapsulated connections and their payloads.

By contrast, the system leverages access control as follows:

- Outer header evaluation—First, prefiltering uses outer headers to handle traffic. You can block or fastpath entire plaintext, passthrough tunnels at this stage.
- Inner header evaluation—Next, the rest of access control (and other features such as QoS) use the innermost detectable level of headers to ensure the most granular level of inspection and handling possible.

  If a passthrough tunnel is not encrypted, the system acts on its individual encapsulated connections at this stage. You must *rezone* a tunnel (see Tunnel Zones and Prefiltering, on page 1928) to act on all its encapsulated connections.

Access control has no insight into encrypted passthrough tunnels. For example, access control rules see a passthrough VPN tunnel as one connection. The system handles the entire tunnel using only the information in its outer, encapsulation header.

Best Practices for Fastpath Prefiltering

When you use the fastpath action in a prefilter rule, the matching traffic bypasses inspection and is simply transmitted through the device. Use this action for traffic that you can trust and that would not benefit from any of the security features available.

The following types of traffic are ideal for fastpathing. For example, you could configure the rules to fastpath any traffic from or to the IP addresses of the endpoints or servers. You can further limit the rule based on ports used.

- VPN traffic that is going through the device. That is, the device is not an endpoint in the VPN topology.
- SQL traffic between trusted endpoints on the internal network.
- Scanner traffic. Scanner probes can create a lot of false-positive responses from intrusion policies.
- Voice/video.
- Backups.
- Management traffic (sftunnel) that traverses Firewall Threat Defense devices. Performing deep inspection on management traffic (using access control policies) can cause issues. You can prefilter based on port TCP/8305 between the management center and managed devices.

Best Practices for Encapsulated Traffic Handling

This topic discusses guidelines for the following types of encapsulated traffic:

- Generic Routing Encapsulation (GRE)
- Point-to-Point Protocol (PPTP)
- IPinIP
- IPv6inIP
- Teredo

GRE Tunnel Limitations

GRE tunnel processing is limited to IPv4 and IPv6 passenger flows. Other protocols, such as PPTP and WCCP, are not supported within the GRE tunnel.

Understand Snort version support for your managed devices

The inspection engine used by managed devices is known as Snort. Snort 3 supports more features than Snort 2. To understand how these affect managed devices on your network, you must know:

- Which versions of Snort your device supports.

Snort version support can be found in the section on bundled components in the *Cisco Firepower Compatibility Guide*.

- How the Firewall Management Center and Firewall Threat Defense software support Snort 2 and Snort 3

  Limitations of Snort 2 and Snort 3 can be found in the *Feature Limitations of Snort 3 for Firewall Management Center-Managed Firewall Threat Defense* topic in the Cisco Secure Firewall Management Center Snort 3 Configuration Guide.

GRE v1 and PPTP bypass outer flow processing

GRE v1 (sometimes referred to as *stateful GRE*) and PPTP traffic bypass outer flow processing.

Passenger flow processing is supported for IPv6inIP and Teredo but the following limitations apply:

- Sessions are over a single tunnel that is not load-balanced
- There is no HA or clustering replication
- Primary and secondary flow relationships are not maintained
- Prefilter policy white and black lists are not supported

GRE v0 sequence number field must be optional

All endpoints sending traffic on the network must send GREv0 traffic with the sequence number field as optional; otherwise, the sequence number field is removed. RFC 1701 and RFC 2784 both specify the sequence field as optional.

How tunnels work with interfaces

Prefilter and access control policy rules are applied to all tunnel types on routed, transparent, inline-set, inline-tap, and passive interfaces.

References

For more information about the GRE and PPTP protocols, see the following:

- RFC 1701, RFC 2784, and RFC 2890 (GRE protocol v0)
- RFC 2637 (PPTP and GRE protocol v1)

Requirements and Prerequisites for Prefilter Policies

Model Support

Threat Defense

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Configure Prefiltering

To perform custom prefiltering, configure prefilter policies and assign the policies to access control policies. It is through the access control policy that prefilter policies get assigned to managed devices.

Only one person should edit a policy at a time, using a single browser window. If multiple users save the same policy, the last saved changes are retained. For your convenience, the system displays information on who (if anyone) is currently editing each policy. To protect the privacy of your session, a warning appears after 30 minutes of inactivity on the policy editor. After 60 minutes, the system discards your changes.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Prefilter**.

Step 2 Click **New Policy** to create a custom prefilter policy.

A new prefilter policy has no rules and a default action of Analyze all tunnel traffic. It performs no logging or tunnel rezoning. You can also **Copy** ( ) or **Edit** ( ) an existing policy.

Step 3 Configure the prefilter policy's default action and its logging options.

- Default action—Choose a default action for supported plaintext, passthrough tunnels: **Analyze all tunnel traffic** (with access control) or **Block all tunnel traffic**.
- Default action logging—Click **Logging** ( ) next to the default action; see *Logging Connections with a Policy Default Action* in the Cisco Secure Firewall Management Center Administration Guide. You can configure default action logging for blocked tunnels only.

Step 4 Configure tunnel and prefilter rules.

In a custom prefilter policy, you can use both kinds of rule, in any order. Create rules depending on the specific type of traffic you want to match and the actions or further analysis you want to perform; see Tunnel vs Prefilter Rules, on page 1916.

Caution
Exercise caution when using tunnel rules to assign tunnel zones. Connections in rezoned tunnels may not match security zone constraints in later evaluation. For more information, see Tunnel Zones and Prefiltering, on page 1928.

For detailed information on configuring rule components, see Tunnel and Prefilter Rule Components, on page 1923.

Step 5 Evaluate rule order. To move a rule, click and drag or use the right-click menu to cut and paste.

Properly creating and ordering rules is a complex task, but one that is essential to building an effective deployment. If you do not plan carefully, rules can preempt other rules or contain invalid configurations. For more information, see Best Practices for Access Control Rules, on page 1738.

Step 6 Save the prefilter policy.

Step 7 For configurations that support tunnel zone constraints, appropriately handle rezoned tunnels.

Match connections in rezoned tunnels by using tunnel zones as source zone constraints.

Step 8 Associate the prefilter policy with the access control policy deployed to your managed devices.

See Associating Other Policies with Access Control, on page 1762.

Step 9 Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Note
When you deploy a prefilter policy, its rules are not applied on the existing tunnel sessions. Hence, traffic on an existing connection is not bound by the new policy that is deployed. In addition, the policy hit count is incremented only for the first packet of a connection that matches a policy. Thus, the traffic on an existing connection that could match a policy is omitted from the hit count. To have the policy rules effectively applied, clear the existing tunnel sessions, and then deploy the policy.

What to do next

If you will deploy time-based rules, specify the time zone of the device to which the policy is assigned. See Time Zone, on page 1003.

Tunnel and Prefilter Rule Components

State (Enabled/Disabled)

By default, rules are enabled. If you disable a rule, the system does not use it and stops generating warnings and errors for that rule.

Position

Rules are numbered, starting at 1. The system matches traffic to rules in top-down order by ascending rule number. The first rule that traffic matches is the rule that handles that traffic, regardless of rule type (tunnel vs prefilter).

Action

A rule's action determines how the system handles and logs matching traffic:

- Fastpath—Exempts matching traffic from all further inspection and control, including access control, identity requirements, and rate limiting. Fastpathing a tunnel fastpaths all encapsulated connections.
- Block—Blocks matching traffic without further inspection of any kind. Blocking a tunnel blocks all encapsulated connections.

- Analyze—Allows traffic to continue to be analyzed by the rest of access control, using inner headers. If passed by access control and any related deep inspection, this traffic may also be rate limited. For tunnel rules, enables rezoning with the Assign Tunnel Zone option.

Note For a prefilter rule with action as **Fastpath** and logging options disabled, you will observe that the end-of-flow events are still generated in the system. The events are not visible on the management center event pages.

Direction (Tunnel Rules Only)

A tunnel rule's direction determines how the system source and destination criteria:

- Match tunnels only from source (unidirectional)—Match source-to-destination traffic only. Matching traffic must originate from one of the specified source interfaces or tunnel endpoints, and leave through one of the destination interfaces or tunnel endpoints. Return traffic for allowed connections is also permitted.
- Match tunnels from source and destination (bidirectional)—Match both source-to-destination traffic and destination-to-source traffic. The effect is identical to writing two unidirectional rules, one the mirror of the other.

Prefilter rules are always unidirectional.

Assign Tunnel Zone (Tunnel Rules Only)

In a tunnel rule, assigning a tunnel zone (whether existing or created on the fly), *rezones* matching tunnels. Rezoning requires the Analyze action.

Rezoning a tunnel allows other configurations—such as access control rules—to recognize all the tunnel's encapsulated connections as belonging together. By using a tunnel's assigned tunnel zone as an interface constraint, you can tailor inspection to its encapsulated connections. For more information, see Tunnel Zones and Prefiltering, on page 1928.

Caution Exercise caution when assigning tunnel zones. Connections in rezoned tunnels may not match security zone constraints in later evaluation. See Using Tunnel Zones, on page 1928 for a brief walkthrough of a tunnel zone implementation, and a discussion of the implications of rezoning without explicitly handling rezoned traffic.

Conditions

Conditions specify the specific traffic the rule handles. Traffic must match all a rule's conditions to match the rule. Each condition type has its own tab in the rule editor.

You can prefilter traffic using the following *outer-header* constraints. You must constrain tunnel rules by encapsulation protocol.

- Interface—Interface Rule Conditions, on page 938
- Network (prefilter rule)/Tunnel Endpoints (tunnel rule)—Network Rule Conditions, on page 939
- VLAN—VLAN Tags Rule Conditions, on page 1787

- Ports (prefilter rule)/Encapsulation and Ports (tunnel rule)—Port Rule Conditions for Prefilter Rules, on page 1926 or Encapsulation Rule Conditions, on page 1927
- Time Range—Time and Day Rule Conditions, on page 1792

Logging

A rule's logging settings govern the records the system keeps of the traffic it handles.

In tunnel and prefilter rules, you can log fastpathed and blocked traffic (the Fastpath and Block actions). For traffic subject to further analysis (the Analyze action), logging in the prefilter policy is disabled, although matching connections may still be logged by other configurations. Logging is performed on inner flows, not on the encapsulating flow. For more information, see *Logging Connections with Tunnel and Prefilter Rules* in the Cisco Secure Firewall Management Center Administration Guide.

Comments

Each time you save changes to a rule you can add comments. For example, you might summarize the overall configuration for the benefit of other users, or note when you change a rule and the reason for the change.

You cannot edit or delete these comments after you save the rule.

Related Topics

Best Practices for Access Control Rules, on page 1738

Prefilter Rule Conditions

Rule conditions enable you to fine-tune your prefilter policy to target the networks you want to control. See one of the following sections for more information.

Interface Rule Conditions

Interface rule conditions control traffic by its source and destination interfaces.

Depending on the rule type and the devices in your deployment, you can use predefined *interface objects* called *security zones* or *interface groups* to build interface conditions. Interface objects segment your network to help you manage and classify traffic flow by grouping interfaces across multiple devices; see Interface, on page 1398. QoS rules can be applied only on routed interfaces.

Tip Constraining rules by interface is one of the best ways to improve system performance. If a rule excludes all of a device's interfaces, that rule does not affect that device's performance.

Network Rule Conditions

Networks control or decrypt traffic by its source and destination IP address, using inner headers. Tunnel rules, which use outer headers, have tunnel endpoint conditions instead of network conditions.

You can use predefined objects to build network conditions, or manually specify individual IP addresses or address blocks.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Note You *cannot* use FDQN network objects in identity rules.

VLAN Tags Rule Conditions

Note VLAN tags in access rules only apply to inline sets. Access rules with VLAN tags do not match traffic on firewall interfaces.

VLAN rule conditions control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic. The system uses the innermost VLAN tag to filter VLAN traffic, with the exception of the prefilter policy, which uses the outermost VLAN tag in its rules.

Note the following Q-in-Q support:

- Firewall Threat Defense on Firepower 4100/9300—Does not support Q-in-Q (supports only one VLAN tag).
- Firewall Threat Defense on all other models:
  - Inline sets and passive interfaces—Supports Q-in-Q, up to 2 VLAN tags.
  - Firewall interfaces—Does not support Q-in-Q (supports only one VLAN tag).

You can use predefined objects to build VLAN conditions, or manually enter any VLAN tag from 1 to 4094. Use a hyphen to specify a range of VLAN tags.

In a cluster, if you encounter problems with VLAN matching, edit the access control policy advanced options, Transport/Network Preprocessor Settings, and select the **Ignore the VLAN header when tracking connections** option.

Port Rule Conditions for Prefilter Rules

Port conditions match traffic based on the source and destination ports. Depending on the rule type, “port” can represent any of the following:

- TCP and UDP—You can control TCP and UDP traffic based on the port. The system represents this configuration using the protocol number in parentheses, plus an optional associated port or port range. For example: TCP(6)/22.
- ICMP—You can control ICMP and ICMPv6 (IPv6-ICMP) traffic based on its internet layer protocol plus an optional type and code. For example: ICMP(1):3:3.
- Protocol—You can control traffic using other protocols that do not use ports.

Using Source and Destination Port Constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as destination port conditions in a single access control rule.

Matching Non-TCP Traffic with Port Conditions

You can match non-port-based protocols. By default, if you do not specify a port condition, you are matching IP traffic. Although you can configure port conditions to match other protocols in prefilter rules, you should use tunnel rules instead when matching GRE, IP in IP, IPv6 in IP, and Toredo Port 3544.

Time and Day Rule Conditions

You can specify a continuous time range or a recurring time period.

For example, a rule can apply only during weekday working hours, or every weekend, or during a holiday shutdown period.

Time-based rules are applied based on the local time of the device that processes the traffic.

Time-based rules are supported only on Firewall Threat Defense devices. If you assign a policy with a time-based rule to a different type of device, the time restriction associated with the rule is ignored on that device. You will see warnings in this case.

Tunnel Rule Conditions

Rule conditions enable you to fine-tune your tunnel policy to target the networks you want to control. For tunnel rules, you can use the following conditions:

- **Interface Objects**—The security zones or interface groups that define the device interfaces through which the connections pass. See Interface Rule Conditions, on page 938.
- **Tunnel Endpoints**—The network objects that define the source and destination IP addresses of the tunnel.
- **VLAN Tags**—The outermost VLAN tag in the tunnel. See VLAN Tags Rule Conditions, on page 1787.
- **Encapsulation and Ports**—The encapsulation protocol of the tunnel. See Encapsulation Rule Conditions, on page 1927.
- **Time Range**—The days and times when the rule is active. If you do not specify a time range, the rule is always active. See Time and Day Rule Conditions, on page 1792.

Encapsulation Rule Conditions

Encapsulation conditions are specific to tunnel rules.

These conditions control certain types of plaintext, passthrough tunnels by their encapsulation protocol. You must choose at least one protocol to match before you can save the rule. You can choose:

- GRE (47)

- IP-in-IP (4)
- IPv6-in-IP (41)
- Teredo (UDP (17)/3455)

Tunnel Zones and Prefiltering

Tunnel zones allow you to use prefiltering to tailor subsequent traffic handling to encapsulated connections.

A special mechanism is required because usually, the system handles traffic using the innermost detectable level of headers. This ensures the most granular level of inspection possible. But it also means that if a passthrough tunnel is not encrypted, the system acts on its individual encapsulated connections; see Passthrough Tunnels and Access Control, on page 1919.

Tunnel zones solve this problem. During the first phase of access control (prefiltering) you can use outer headers to identify certain types of plaintext, passthrough tunnels. Then, you can *rezone* those tunnels by assigning a custom *tunnel zone*.

Rezoning a tunnel allows other configurations—such as access control rules—to recognize all the tunnel's encapsulated connections as belonging together. By using a tunnel's assigned tunnel zone as an interface constraint, you can tailor inspection to its encapsulated connections.

Despite its name, a tunnel zone is not a security zone. A tunnel zone does not represent a set of interfaces. It is more accurate to think of a tunnel zone as a tag that, in some cases, replaces the security zone associated with an encapsulated connection.

Caution For configurations that support tunnel zone constraints, connections in rezoned tunnels do **not** match security zone constraints. For example, after you rezone a tunnel, access control rules can match its encapsulated connections with their newly assigned *tunnel* zone, but not with any original *security* zone.

See Using Tunnel Zones, on page 1928 for a brief walkthrough of a tunnel zone implementation, and a discussion of the implications of rezoning without explicitly handling rezoned traffic.

Configurations Supporting Tunnel Zone Constraints

Only access control rules support tunnel zone constraints.

No other configurations support tunnel zone constraints. For example, you cannot use QoS to rate limit a plaintext tunnel as a whole; you can only rate limit its individual encapsulated sessions.

Using Tunnel Zones

This example procedure summarizes how you might rezone GRE tunnels for further analysis, using tunnel zones. You can adapt the concepts described in this example to other scenarios where you need to tailor traffic inspection to connections encapsulated in plaintext, passthrough tunnels.

Consider a situation where your organization's internal traffic flows through the Trusted security zone. The Trusted security zone represents a set of interfaces across multiple managed devices deployed in various locations. Your organization's security policy requires that you allow internal traffic after deep inspection for exploits and malware.

Internal traffic sometimes includes plaintext, passthrough, GRE tunnels between particular endpoints. Because the traffic profile of this encapsulated traffic is different from your "normal" interoffice activity—perhaps it is known and benign—you can limit inspection of certain encapsulated connections while still complying with your security policy.

In this example, after you deploy configuration changes:

- Plaintext, passthrough, GRE-encapsulated tunnels detected in the Trusted zone have their individual encapsulated connections evaluated by one set of intrusion and file policies.
- All other traffic in the Trusted zone is evaluated with a different set of intrusion and file policies.

You accomplish this task by *rezoning* GRE tunnels. Rezoning ensures that access control associates GRE-encapsulated connections with a custom *tunnel* zone, rather than their original Trusted *security* zone. Rezoning is required due to the way access control handles encapsulated traffic; see Passthrough Tunnels and Access Control, on page 1919 and Tunnel Zones and Prefiltering, on page 1928.

Procedure

Step 1 Configure custom intrusion and file policies that tailor deep inspection to encapsulated traffic, and another set of intrusion and file policies tailored to nonencapsulated traffic.

Step 2 Configure custom prefiltering to rezone GRE tunnels flowing through the Trusted security zone.

Create a custom prefilter policy and associate it with access control. In that custom prefilter policy, create a tunnel rule (in this example, **`GRE_tunnel_rezone`**) and a corresponding tunnel zone (**`GRE_tunnel`**). For more information, see Configure Prefiltering, on page 1922.

Table 103: GRE_tunnel_rezone Tunnel Rule

Rule Component	Description
Interface object condition	Match internal-only tunnels by using the Trusted security zone as both a Source Interface Object and Destination Interface Object constraint.
Tunnel endpoint condition	Specify the source and destination endpoints for the GRE tunnels used in your organization. Tunnel rules are bidirectional by default. If you do not change the **Match tunnels from...** option, it does not matter which endpoints you specify as source and which as destination.
Encapsulation condition	Match GRE traffic.
Assign Tunnel Zone	Create the **`GRE_tunnel`** tunnel zone, and assign it to tunnels that match the rule.
Action	Analyze (with the rest of access control).

Step 3 Configure access control to handle connections in rezoned tunnels.

In the access control policy deployed to your managed devices, configure a rule (in this example, **`GRE_inspection`**) that handles the traffic you rezoned. For more information, see Create and Edit Access Control Rules, on page 1783.

Table 104: GRE_inspection Access Control Rule

Rule Component	Description
Security zone condition	Match rezoned tunnels by using the GRE_tunnel security zone as a Source Zone constraint.
Action	Allow, with deep inspection enabled. Choose the file and intrusion policies tailored to inspect encapsulated internal traffic.

Caution
If you skip this step, the rezoned connections may match **any** access control rule not constrained by security zone. If the rezoned connections do not match any access control rules, they are handled by the access control policy default action. Make sure this is your intent.

Step 4 Configure access control to handle nonencapsulated connections flowing through the Trusted security zone.

In the same access control policy, configure a rule (in this example, **`internal_default_inspection`**) that handles non-rezoned traffic in the Trusted security zone.

Table 105: internal_default_inspection Access Control Rule

Rule Component	Description
Security zone condition	Match non-rezoned internal-only traffic by using the Trusted security zone as both a Source Zone and Destination Zone constraint.
Action	Allow, with deep inspection enabled. Choose the file and intrusion policies tailored to inspect nonencapsulated internal traffic.

Step 5 Evaluate the position of the new access control rules relative to preexisting rules. Change rule order if necessary.

If you place the two new access control rules next to each other, it does not matter which you place first. Because you rezoned GRE tunnels, the two rules cannot preempt each other.

Step 6 Save all changed configurations.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Creating Tunnel Zones

The following procedure explains how to create a tunnel zone in the object manager. You can also create zones when editing a tunnel rule.

Procedure

Step 1 Choose **Objects** > **Object Management**.

Step 2 Choose **Tunnel Zone** from the list of object types.

Step 3 Click **Add Tunnel Zone**.

Step 4 Enter a **Name** and, optionally, a **Description**.

Step 5 Click **Save**.

What to do next

- Assign tunnel zones to plaintext, passthrough tunnels as part of custom prefiltering; see Configure Prefiltering, on page 1922.

Moving Prefilter Rules to an Access Control Policy

You can move prefilter rules from a prefilter policy to the associated access control policy.

Before you begin

Note the following conditions before you proceed:

- Only prefilter rules can be moved to an access control policy. Tunnel rules cannot be moved.
- The prefilter rules can be moved only to the associated access control policy.
- The prefilter rules with configured interface groups cannot be moved.
- The **Action** parameter in the prefilter rule is changed to a suitable action in the access control rule when moved. To know what each action in the prefilter rule maps to, see the following table:

Action in the prefilter rule	Action in the access control rule
Analyze	Allow
Block	Block
Fastpath	Trust

- Similarly, based on the action configured in the prefilter rule, the logging configuration is set to an appropriate setting after the rule is moved, as mentioned in the following table.

Action in the prefilter rule	Enabled Logging configurations in the access control rule
Analyze	None of the log settings are enabled.

Action in the prefilter rule	Enabled Logging configurations in the access control rule
Block	• Log at Beginning of Connection • Event Viewer • Syslog Server • SNMP Trap
Fastpath	• Log at Beginning of Connection • Log at End of Connection • Event Viewer • Syslog Server • SNMP Trap

- The comments in the prefilter rule configuration are lost after moving the rule. However, a new comment is added in the moved rule mentioning the source prefilter policy.
- While moving rules from the source policy, if another user modifies those rules, the Firewall Management Center displays a message. You may continue with the process after refreshing the page.

Procedure

Step 1 In the prefilter policy editor, select the rules that you want to move with a left-click on your mouse.

Tip
To select multiple rules, use the Ctrl (Control) key on your keyboard.

Step 2 Right-click the selected rules and choose **Move to another policy**.

Step 3 Select the destination access control policy from the **Access Policy** drop-down list.

Step 4 From the **Place Rules** drop-down list, choose where you want to position the moved rules:

- To position as the last set of rules in the **Default** section, choose **At the bottom (within the Default section)**.
- To position as the first set of rules in the **Mandatory** section, choose **At the top (within the Mandatory section)**.

Step 5 Click **Move**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Prefilter Policy Hit Counts

Hit count indicates the number of times a policy rule has triggered for a matching connection.

For complete information on viewing prefilter policy hit counts, see Viewing Rule Hit Counts, on page 1764.

Large Flow Offloads

On the 3100/, 4100/4200/, 9300, certain traffic that you configure to be fastpathed by a prefilter policy is handled by the hardware (specifically, in the NIC), not by your Firewall Threat Defense software. Offloading these connection flows results in higher throughput and lower latency, especially for data-intensive applications such as large file transfers. This feature is especially useful for data centers. This is called *static flow offload*.

In addition, by default, Firewall Threat Defense devices offload flows based on other criteria, including trust. This is called *dynamic flow offload*.

Offloaded flows continue to receive limited stateful inspection, such as basic TCP flag and option checking. The system can selectively escalate packets to the firewall system for further processing if necessary.

Examples of applications that can benefit from offloading large flows are:

- High Performance Computing (HPC) Research sites, where the Firewall Threat Defense device is deployed between storage and high compute stations. When one research site backs up using FTP file transfer or file sync over NFS, the large amount of data traffic affects all connections. Offloading FTP file transfer and file sync over NFS reduces the impact on other traffic.
- High Frequency Trading (HFT), where the Firewall Threat Defense device is deployed between workstations and the Exchange, mainly for compliance purposes. Security is usually not a concern, but latency is a major concern.

The following flows can be offloaded:

- (Static flow offload only.) Connections that are fastpathed by the prefilter policy.
- Standard or 802.1Q tagged Ethernet frames only.
- (Dynamic flow offload only):
  - Inspected flows that the inspection engine decides no longer need inspection. These flows include:
    - Flows handled by access control rules that apply the Trust action and are based on security zone, source and destination network and port matching only.
    - TLS/SSL flows that are not selected for decryption using a decryption policy.
    - Flows that are trusted by the Intelligent Application Bypass (IAB) policy either explicitly or due to exceeding flow bypass thresholds.
    - Flows that match file or intrusion policies that result in trusting the flow.
    - Any allowed flow that no longer needs to be inspected.
  - The following IPS preprocessor inspected flows:
    - SSH and SMTP.

- FTP preprocessor secondary connections.
- Session Initiation Protocol (SIP) preprocessor secondary connections.

- Intrusion rules that use keywords (also referred to as *options*)

Important For details, exceptions, and limitations to the above, see Flow Offload Limitations, on page 1934.

Use Static Flow Offload

To offload eligible traffic to hardware, create a prefilter policy rule that applies the **Fastpath** action. Use prefilter rules for TCP/UDP, and tunnel rules for GRE.

(Not recommended.) To disable static flow offload and as a by-product, dynamic flow-offload, use FlexConfig to run the **no flow-offload enable** command. For information about this command, see the *Cisco ASA Series Command Reference*, available from https://www.cisco.com/c/en/us/support/security/adaptive-security-appliance-asa-software/products-command-reference-list.html.

Use Dynamic Flow Offload

Dynamic flow offload is enabled by default.

To disable dynamic offload:

configure flow-offload dynamic whitelist disable

To re-enable dynamic offload:

configure flow-offload dynamic whitelist enable

Note that dynamic offload occurs only if static flow offload is enabled, regardless of whether prefiltering is configured.

Flow Offload Limitations

Not all flows can be offloaded. Even after offload, a flow can be removed from being offloaded under certain conditions. Following are some of the limitations:

Device Limitations

The feature is supported on the following devices:

- Secure Firewall 3100
- Secure Firewall 4200
- Firepower 4100/9300

Flows that cannot be offloaded

The following types of flows cannot be offloaded.

- Any flows that do not use IPv4 addressing, such as IPv6 addressing.
- Flows for any protocol other than TCP, UDP, and GRE.

Note PPTP GRE connections cannot be offloaded.

- Flows on interfaces configured in passive, inline, or inline tap mode. Routed and switched interfaces are the only types supported.
- (Secure Firewall 3100/4200.) Offload based on inner header for tunnelled flows.
- (Secure Firewall 3100/4200.) Multi-instance offload.
- Flows that require inspection by Snort or other inspection engines. In some cases, such as FTP, the secondary data channel can be offloaded although the control channel cannot be offloaded.
- IPsec and TLS/DTLS VPN connections that terminate on the device.
- Flows that require encryption or decryption. For example, connections decrypted due to a decryption policy.
- Multicast flows in routed mode. They are supported in transparent mode if there are only two member interfaces in a bridge group.
- TCP Intercept flows.
- TCP state bypass flows. You cannot configure flow offload and TCP state bypass on the same traffic.
- Flows tagged with security groups.
- Reverse flows that are forwarded from a different cluster node, in the case of asymmetric flows in a cluster.
- Centralized flows in a cluster, if the flow owner is not the control unit.
- Flows that include IP options cannot be dynamically offloaded.

Additional Limitations

- Flow offload and Dead Connection Detection (DCD) are not compatible. Do not configure DCD on connections that can be offloaded.
- On Firepower 4100 and 9300, If more than one flow that matches flow offload conditions are queued to be offloaded at the same time to the same location on the hardware, only the first flow is offloaded. The other flows are processed normally. This is called a *collision*. Use the **show flow-offload flow** command in the CLI to display statistics for this situation.

  On Secure Firewall 3100 and 4200, the cuckoo hashing technique is used to resolve hash collisions internally in a more efficient way due to which the **show flow-offload flow** command will not show any increments in collision counter statistics.
- Dynamic flow offload disables all TCP normalizer checks.
- Although offloaded flows pass through FXOS interfaces, statistics for these flows do not appear on the logical device interface. Thus, logical device interface counters and packet rates do not reflect offloaded flows.

Conditions for reversing offload

After a flow is offloaded, packets within the flow are returned to the Firewall Threat Defense for further processing if they meet the following conditions:

- They include TCP options other than Timestamp.
- They are fragmented.
- They are subject to Equal-Cost Multi-Path (ECMP) routing, and ingress packets move from one interface to another.

History for Prefiltering

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Moving prefilter rules to an access control policy	6.7	Any	You can move prefilter rules from a prefilter policy to the associated access control policy. New/modified pages: In the prefilter policy page, the right-click menu for the selected rules provides a new **Move to another policy** option. Supported platforms: Firewall Management Center
Time-based rules	6.6	Any	Ability to apply prefilter and tunnel rules depending on the date and time, as determined by the time zone of the Firewall Threat Defense device. See description in History for Access Control Rules, on page 1801.
View Object Details from prefilter rule page	6.6	Any	Feature introduced: Option to view details for an object or object group when viewing prefilter rules. New options: Right-clicking a value in any of the following columns in the prefilter rule list offers an option to view object details: Source Networks, Destination Networks, Source Port, Destination Port, and VLAN Tag. Supported platforms: Secure Firewall Management Center

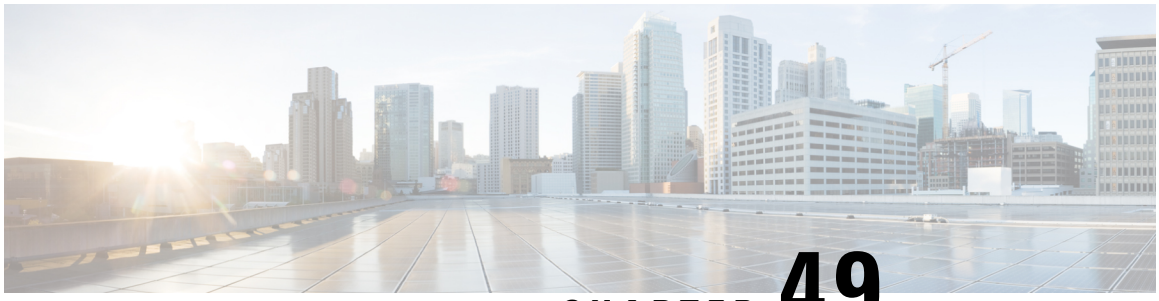

CHAPTER 49

Service Policies

You can use Threat Defense Service Policies to apply services to specific traffic classes. For example, you can use a service policy to create a timeout configuration that is specific to a particular TCP application, as opposed to one that applies to all TCP applications. A service policy consists of multiple actions or rules applied to an interface or applied globally.

- About Threat Defense Service Policies, on page 1937
- Requirements and Prerequisites for Service Policies, on page 1939
- Guidelines and Limitations for Service Policies, on page 1939
- Configure Threat Defense Service Policies, on page 1940
- Examples for Service Policy Rules, on page 1948
- Monitoring Service Policies, on page 1953
- History for Threat Defense Service Policy, on page 1954

About Threat Defense Service Policies

You can use Threat Defense Service Policies to apply services to specific traffic classes. With service policies, you are not limited to applying the same services to all connections that enter the device or a given interface.

A traffic class is a combination of the interface and an extended access control list (ACL). The ACL “allow” rules determine which connections are part of the class. Any “denied” traffic in the ACL simply does not have the service applied to it: these connections are not actually dropped. You can use IP addresses and TCP/UCP ports to identify matching connections as precisely as you require.

There are two types of traffic class:

- Interface-based rules—If you specify a security zone or interface group in a service policy rule, the rule applies to the ACL “allowed” traffic that goes through any interface that is part of the interface objects.

  For a given feature, interface-based rules applied to the ingress interface always take precedence over global rules: if an ingress interface-based rule applies to a connection, any matching global rule is ignored. If no ingress interface or global rule applies, then an interface service rule on the egress interface is applied.

- Global rules—These rules apply to all interfaces. If an interface-based rule does not apply to a connection, the global rules are checked and applied to any connections that the ACL “allows.” If none apply, then the connections proceed without any services applied.

A given connection can match only one traffic class, either interface-based or global, for a given feature. There should be at most one rule for a given interface object/traffic flow combination.

Service policy rules are applied after access control rules. These services are configured only for connections you are allowing.

How Service Policies Relate to FlexConfig and Other Features

Prior to version 6.3(0), you could configure connection-related service rules using the TCP_Embryonic_Conn_Limit and TCP_Embryonic_Conn_Timeout pre-defined FlexConfig objects. You should remove those objects and redo your rules using the Threat Defense Service Policy. If you created any custom FlexConfig objects to implement any of these connection-related features (that is, **set connection** commands), you should also remove those objects and implement the features through the service policy.

Because connection-related service policy features are treated as a separate feature group from other service-rule implemented features, you should not run into problems with overlapping traffic classes. However, please be mindful when configuring the following:

- QoS Policy rules are implemented using the service policy CLI. These rules are applied before connection-based service policy rules. However, both QoS and connection settings can be applied to the same or overlapping traffic classes.
- You can use FlexConfig policies to implement customized application inspections and NetFlow. Use the **show running-config** command to examine the CLI that already configures service rules, including the **policy-map**, **class-map**, and **service-policy** commands. Netflow and application inspection are compatible with QoS and connection settings, but you need to understand the existing configuration before implementing FlexConfig. Connection settings are applied before application inspections and Netflow.

Note Traffic classes that are created from the Threat Defense Service Policy are named **class_map_***ACLname*, where *ACLname* is the name of the extended ACL object used in the service policy rule.

What Are Connection Settings?

Connection settings comprise a variety of features related to managing traffic connections, such as a TCP flow through the Firewall Threat Defense. Some features are named components that you would configure to supply specific services.

Connection settings include the following:

- **Global timeouts for various protocols**—All global timeouts have default values, so you need to change them only if you are experiencing premature connection loss. You configure global timeouts in the Firepower Threat Defense Platform policy. Select **Devices** > **Platform Settings**.
- **Connection timeouts per traffic class**—You can override the global timeouts for specific types of traffic using service policies. All traffic class timeouts have default values, so you do not have to set them.
- **Connection limits and TCP Intercept**—By default, there are no limits on how many connections can go through (or to) the Firewall Threat Defense. You can set limits on particular traffic classes using service policy rules to protect servers from denial of service (DoS) attacks. Particularly, you can set limits on embryonic connections (those that have not finished the TCP handshake), which protects against SYN

flooding attacks. When embryonic limits are exceeded, the TCP Intercept component gets involved to proxy connections and ensure that attacks are throttled.

- **Dead Connection Detection (DCD)**—If you have persistent connections that are valid but often idle, so that they get closed because they exceed idle timeout settings, you can enable Dead Connection Detection to identify idle but valid connections and keep them alive (by resetting their idle timers). Whenever idle times are exceeded, DCD probes both sides of the connection to see if both sides agree the connection is valid. The **show service-policy** command output includes counters to show the amount of activity from DCD. You can use the **show conn detail** command to get information about the initiator and responder and how often each has sent probes.
- **TCP sequence randomization**—Each TCP connection has two initial sequence numbers (ISN): one generated by the client and one generated by the server. By default, the Firewall Threat Defense randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions. Randomization prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session. However, TCP sequence randomization effectively breaks TCP SACK (Selective Acknowledgement), as the sequence numbers the client sees are different from what the server sees. You can disable randomization per traffic class if desired.
- **TCP Normalization**—The TCP Normalizer protects against abnormal packets. You can configure how some types of packet abnormalities are handled by traffic class. You can configure TCP Normalization using the FlexConfig policy.
- **TCP State Bypass**—You can bypass TCP state checking if you use asymmetrical routing in your network.

Requirements and Prerequisites for Service Policies

Model Support

Threat Defense

Supported Domains

Any

User Roles

Admin

Access Admin

Network Admin

Guidelines and Limitations for Service Policies

- Service policies apply to routed or switch interfaces only, in either routed or transparent mode. They do not apply to inline set or passive interfaces.
- You can have at most 25 traffic classes for a given interface or the global policy. Specifically, this means that you cannot have more than 25 service policy rules for the global policy for a given security zone or interface group. However, for interfaces, because the same interface can appear in both a security zone

and interface group, be aware that the actual limitation is based on the interfaces, and not the zone/group. Thus, you might be prevented from having 25 rules per zone/group based on the membership of your zones/groups.

- You can have at most one rule for a given interface object/traffic flow combination.
- When you make service policy changes to the configuration, all new connections use the new service policy. Existing connections continue to use the policy that was configured at the time of the connection establishment. If you want all connections to immediately use the new policy, you need to disconnect the current connections so they can reconnect using the new policy. From an SSH or Console CLI session, enter the **clear conn** or **clear local-host** commands.

Configure Threat Defense Service Policies

You can use Threat Defense Service Policies to apply services to specific traffic classes. For example, you can use a service policy to create a timeout configuration that is specific to a particular TCP application, as opposed to one that applies to all TCP applications. A service policy consists of multiple actions or rules applied to an interface or applied globally.

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**, and click **Edit** (✎) for the access control policy whose Threat Defense Service Policy you want to edit.

Step 2 Click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 3 Click **Edit** (✎) in the **Threat Defense Service Policy** group.

A dialog box opens that shows the existing policy. The policy consists of an ordered list of rules, separated between global rules (which apply to all interfaces) and interface-based rules. The table shows the interface object and extended access control list name (which combined defines the traffic class for the rule), and the services applied.

Step 4 Do any of the following:

- Click **Add Rule** to create a new rule. See Configure a Service Policy Rule, on page 1941.
- Click **Edit** (✎) to edit an existing rule. See Configure a Service Policy Rule, on page 1941.
- Click **Delete** (🗑) to delete a rule.
- Click a rule and drag it to a new location to move it. You cannot drag rules between the interface and global lists, instead you must edit the rule to change the interface/global setting. The first rule in the list that matches a connection is applied to the connection.

Step 5 Click **OK** when you are finished editing the policy.

Step 6 Click **Save** on **Advanced** window. The changes are not saved until you click save.

Configure a Service Policy Rule

Configure service policy rules to apply services to specific traffic classes.

Before you begin

Go to **Objects** > **Object Management** > **Access List** > **Extended** and create an the extended access list that defines the traffic to which the rule applies. The rule is applied to any connections that match Allow rules in the extended access list. Define the ACL rules precisely, so that your service policy rule applies to only the traffic that requires the service.

If you are creating an interface-based rule, you must also have configured the interfaces on the assigned devices and added them to security zones or interface groups.

Procedure

Step 1 If you are not already in the Threat Defense Service Policy dialog box, choose **Policies** > **Access Control heading** > **Access Control**, edit the access control policy, select **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line, then edit the **Threat Defense Service Policy**.

Step 2 Do any of the following:

- Click **Add Rule** to create a new rule.
- Click **Edit** (✎) to edit an existing rule.

The service policy rule wizard opens to step you through the process of configuring the rule.

Step 3 In the **Interface Object** step, select the option that defines the interfaces that will use the policy.

- **Apply Globally**—Select this option to create a global rule, which applies to all interfaces.
- **Select Interface Objects**—Select this option to create an interface-based rule. Then, select the security zones or interface objects that contain the desired interfaces, and click **>** to move them to the **Next** selected list. The service policy rule will be configured on each interface contained in the selected objects; it is not configured on the zone/group itself.

Click when the interface criteria is complete.

Step 4 In the **Traffic Flow** step, select the extended ACL object that defines the connections to which the rule applies, then click **Next**.

Step 5 In the **Connection Setting** step, configure the services to apply to this traffic class.

- **Enable TCP State Bypass** (TCP connections only)—Implement TCP State Bypass. Connections subject to TCP State Bypass are not inspected by any inspection engines, and they bypass all TCP state checking and TCP normalization. For detailed information, see Bypass TCP State Checks for Asymetrical Routing (TCP State Bypass), on page 1943.

  Note
  Use TCP State Bypass for troubleshooting purposes or when asymmetric routing cannot be resolved. This feature disables multiple security features, which can cause a high number of connections if you do not implement it properly with a narrowly-defined traffic class.

- **Randomize TCP Sequence Number** (TCP connections only)—Whether to enable or disable TCP sequence number randomization. Randomization is enabled by default. For more information, see Disable TCP Sequence Randomization, on page 1947.
- **Enable Decrement TTL** (TCP connections only)—Decrement the time-to-live (TTL) on packets that match the class. If you decrement time to live, packets with a TTL of 1 will be dropped, but a connection will be opened for the session on the assumption that the connection might contain packets with a greater TTL. Note that some packets, such as OSPF hello packets, are sent with TTL = 1, so decrementing time to live can have unexpected consequences.

  Note
  If you want the Firewall Threat Defense device to appear on traceroutes, you must configure the decrement TTL option and also set the ICMP unreachable rate limit in the platform settings policy. See Make the Firewall Threat Defense Device Appear on Traceroutes, on page 1951.

- **Connections**—Limits for the number of connections allowed for the entire class. You can configure these options:
  - **Maximum TCP and UDP** (TCP/UDP connections only)—The maximum number of simultaneous connections that are allowed, between 0 and 2000000, for the entire class. For TCP, this count applies to established connections only. The default is 0, which allows unlimited connections. Because the limit is applied to a class, one attacking host can consume all the connections and leave none for the rest of the hosts that are matched to the class. Set the per-client limit to ameliorate this problem.
  - **Maximum Embryonic** (TCP connections only)—The maximum number of simultaneous embryonic TCP connections (those that have not finished the TCP handshake) allowed, between 0 and 2000000. The default is 0, which allows unlimited connections. By setting a non-zero limit, you enable TCP Intercept, which protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. Also set the per-client options to protect against SYN flooding. For more information, see Protect Servers from a SYN Flood DoS Attack (TCP Intercept), on page 1948.
- **Connections Per Client**—Limits for the number of connections allowed for a given client (source IP address). You can configure these options:
  - **Maximum TCP and UDP** (TCP/UDP connections only)—The maximum number of simultaneous connections allowed per client, between 0 and 2000000. For TCP, this includes established, half-open (embryonic), and half-closed connections. The default is 0, which allows unlimited connections. This option restricts the maximum number of simultaneous connections that are allowed for each host that is matched to the class.
  - **Maximum Embryonic** (TCP connections only)—The maximum number of simultaneous embryonic TCP connections allowed per client, between 0 and 2000000. The default is 0, which allows unlimited connections. For more information, see Protect Servers from a SYN Flood DoS Attack (TCP Intercept), on page 1948.
- **Connections Syn Cookie MSS**—The server maximum segment size (MSS) for SYN-cookie generation for embryonic connections upon reaching the embryonic connections limit, from 48 to 65535 . The default is 1380. This setting is meaningful only if you configure **Maximum Embryonic** for connections or per-client, or both.
- **Connections Timeout**—The timeout settings to apply to the traffic class. These timeouts override the global timeouts defined in the platform settings policy. You can configure the following:

- **Embryonic** (TCP connections only)—The timeout period until a TCP embryonic (half-open) connection is closed, between 0:0:5 and 1193:00:00. The default is 0:0:30.
- **Half Closed** (TCP connections only)—The idle timeout period until a half-closed connection is closed, between 0:0:30 and 1193:0:0. The default is 0:10:0. Half-closed connections are not affected by Dead Connection Detection (DCD). Also, the system does not send a reset when taking down half-closed connections.
- **Idle** (TCP, UDP, ICMP, IP connections)—The idle timeout period after which an established connection of any protocol closes, between 0:0:1 and 1193:0:0. The default is 1:0:0, unless you select the TCP State Bypass option, where the default is 0:2:0.
- **Reset Connection Upon Timeout** (TCP connections only)—Whether to send a TCP RST packet to both end systems after idle connections are removed.

- **Detect Dead Connections** (TCP connections only)—Whether to enable Dead Connection Detection (DCD). Before expiring an idle connection, the system probes the end hosts to determine if the connection is valid. If both hosts respond, the connection is preserved, otherwise the connection is freed. When operating in transparent firewall mode, you must configure static routes for the endpoints. You cannot configure DCD on connections that are also offloaded, so do not configure DCD on connections you are fast-pathing in the prefilter policy. Use the **show conn detail** command in the Firewall Threat Defense CLI to track how many DCD probes have been sent by the initiator and responder.

  Configure the following options:

  - **Detection Timeout**—The time duration in hh:mm:ss format to wait after each unresponsive DCD probe before sending another probe, between 0:0:1 and 24:0:0. The default is 0:0:15.

    For systems that are operating in a cluster or high-availability configuration, we recommend that you do not set the interval to less than one minute (0:1:0). If the connection needs to be moved between systems, the changes required take longer than 30 seconds, and the connection might be deleted before the change is accomplished.
  - **Detection Retries**—The number of consecutive failed retries for DCD before declaring the connection dead, from 1 to 255. The default is 5.

Step 6 Click **Finish** to save your changes.

The rule is added to the bottom of the appropriate list, either Interfaces or Global. Global rules are matched in top-down order. Rules in the Interfaces list are matched in top down order for each interface object. Place rules for narrowly-defined traffic classes above broader rules, to ensure the right services get applied. You can move rules within each list by using drag and drop. You cannot move rules between lists.

Bypass TCP State Checks for Asymetrical Routing (TCP State Bypass)

If you have an asymetrical routing environment in your network, where the outbound and inbound flow for a given connection can go through two different Firewall Threat Defense devices, you need to implement TCP State Bypass on the affected traffic.

However, TCP State Bypass weakens the security of your network, so you should apply bypass on very specific, limited traffic classes.

The following topics explain the problem and solution in more detail.

The Asymetrical Routing Problem

By default, all traffic that goes through the Firewall Threat Defense is inspected using the Adaptive Security Algorithm and is either allowed through or dropped based on the security policy. The Firewall Threat Defense maximizes the firewall performance by checking the state of each packet (new connection or established connection) and assigning it to either the session management path (a new connection SYN packet), the fast path (an established connection), or the control plane path (advanced inspection).

TCP packets that match existing connections in the fast path can pass through the Firewall Threat Defense without rechecking every aspect of the security policy. This feature maximizes performance. However, the method of establishing the session in the fast path using the SYN packet, and the checks that occur in the fast path (such as TCP sequence number), can stand in the way of asymmetrical routing solutions: both the outbound and inbound flow of a connection must pass through the same Firewall Threat Defense device.

For example, a new connection goes to Security Appliance 1. The SYN packet goes through the session management path, and an entry for the connection is added to the fast path table. If subsequent packets of this connection go through Security Appliance 1, then the packets match the entry in the fast path, and are passed through. But if subsequent packets go to Security Appliance 2, where there was not a SYN packet that went through the session management path, then there is no entry in the fast path for the connection, and the packets are dropped. The following figure shows an asymmetric routing example where the outbound traffic goes through a different Firewall Threat Defense than the inbound traffic:

Figure 440: Asymmetric Routing

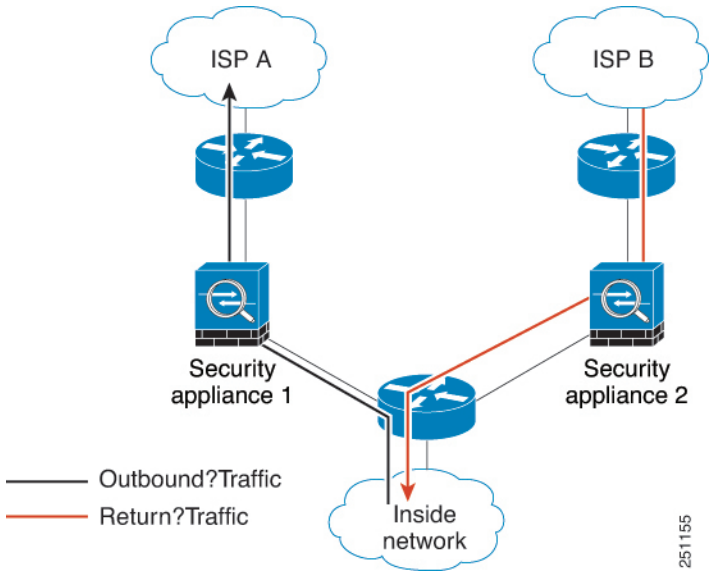

If you have asymmetric routing configured on upstream routers, and traffic alternates between two Firewall Threat Defense devices, then you can configure TCP state bypass for specific traffic. TCP state bypass alters the way sessions are established in the fast path and disables the fast path checks. This feature treats TCP traffic much as it treats a UDP connection: when a non-SYN packet matching the specified networks enters the Firewall Threat Defense device, and there is not a fast path entry, then the packet goes through the session management path to establish the connection in the fast path. Once in the fast path, the traffic bypasses the fast path checks.

Guidelines and Limitations for TCP State Bypass

TCP State Bypass Unsupported Features

The following features are not supported when you use TCP state bypass:

- Application inspection—Inspection requires both inbound and outbound traffic to go through the same Firewall Threat Defense, so inspection is not applied to TCP state bypass traffic.
- Snort inspection—Inspection requires both inbound and outbound traffic to go through the same device. However, Snort inspection is not automatically bypassed for TCP state bypass traffic. You must also configure a prefilter fastpath rule for the same traffic class for which you configure TCP state bypass. Otherwise, packets might be dropped unexpectedly because the TCP Normalizer is also not engaged.
- TCP Intercept, maximum embryonic connection limit, TCP sequence number randomization—The Firewall Threat Defense does not keep track of the state of the connection, so these features are not applied.
- TCP normalization—The TCP normalizer is disabled.
- Stateful failover.
- TLS server identity discovery cannot be used with TCP state bypass on an inline or inline tap interface.

TCP State Bypass NAT Guidelines

Because the translation session is established separately for each Firewall Threat Defense, be sure to configure static NAT on both devices for TCP state bypass traffic. If you use dynamic NAT, the address chosen for the session on Device 1 will differ from the address chosen for the session on Device 2.

Configure TCP State Bypass

To bypass TCP state checking in asymetrical routing environments, carefully define a traffic class that applies to the affected hosts or networks only, then enable TCP State Bypass on the traffic class using a service policy. You must also configure a corresponding prefilter fastpath policy for the same traffic to ensure the traffic also bypasses inspection.

Because bypass reduces the security of the network, limit its application as much as possible.

Procedure

Step 1 Create the extended ACL that defines the traffic class.

For example, to define a traffic class for TCP traffic from 10.1.1.1 to 10.2.2.2, do the following:

a) Choose **Objects** > **Object Management**.
b) Choose **Access List** > **Extended** from the table of contents.
c) Click **Add Extended Access List**.
d) Enter a **Name** for the object, for example, bypass.
e) Click **Add** to add a rule.
f) Keep **Allow** for the action.

g) Enter 10.1.1.1 beneath the **Source** list and click **Add**, and 10.2.2.2 beneath the **Destination** list, and click **Add**.
h) Click **Port**, select **TCP (6)** beneath the **Selected Source Ports** list, and click **Add**. Do not enter a port number, simply add TCP as the protocol, which will cover all ports.
i) Click **Add** on the Extended Access List Entry dialog box to add the rule to the ACL.
j) Click **Save** on the Extended Access List Object dialog box to save the ACL object.

Step 2 Configure the TCP state bypass service policy rule.

For example, to configure TCP state bypass for this traffic class globally, do the following:

a) Choose **Policies** > **Access Control heading** > **Access Control**, and edit the policy assigned to the devices that require this service.
b) Click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line, and click **Edit** (✎) for the **Threat Defense Service Policy**.
c) Click **Add Rule**.
d) Select **Apply Globally** > **Next**.
e) Select the extended ACL object you created for this rule and click **Next**.
f) Select **Enable TCP State Bypass**.
g) (Optional.) Adjust the **Idle** timeout for bypassed connections. The default is 2 minutes.
h) Click **Finish** to add the rule. If necessary, drag and drop the rule to the desired position in the service policy.
i) Click **OK** to save the changes to the service policy.
j) Click **Save** on **Advanced** to save the changes to the access control policy.

Step 3 Configure a prefilter fastpath rule for the traffic class.

You cannot use the ACL object in the prefilter rule, so you need to recreate the traffic class either directly in the prefilter rule, or by first creating network objects that define the class.

The following procedure assumes that you already have a prefilter policy attached to the access control policy. If you have not created a prefilter policy yet, go to **Policies** > **Prefilter** and first create the policy. You can then follow this procedure to attach it to the access control policy and create the rule.

Keeping with our example, this procedure creates a fastpath rule for TCP traffic from 10.1.1.1 to 10.2.2.2.

a) Choose **Policies** > **Access Control heading** > **Access Control**, and edit the policy that has the TCP bypass service policy rule.
b) Click the link for the **Prefilter Policy**, which is to the left immediately under the policy description.
c) In the Prefilter Policy dialog box, select the policy to assign to the device if the correct one is not already selected. Do not click OK yet.

   Because you cannot add rules to the default prefilter policy, you must choose a custom policy.

d) In the Prefilter Policy dialog box, click the **Edit** (✎). This action opens a new browser window where you can edit the policy.
e) Click **Add Prefilter Rule** and configure a rule with the following properties.

   - **Name**—Any name that you fined meaningful will do, such as TCPBypass.
   - **Action**—Select **Fastpath**.
   - **Interface Objects**—If you configured TCP state bypass as a global rule, leave the default, any, for both source and destination. If you created an interface-based rule, select the same interface objects you used for rule in the **Source Interface Objects** list, and keep any as the destination.

- **Networks**—Add 10.1.1.1 to the **Source Networks** list, and 10.2.2.2 to the **Destination Networks** list. You can either use network objects or manually add the addresses.
- **Ports**—Under **Selected Source Ports**, select TCP(6), **do not enter a port**, and click **Add**. This will apply the rule to all (and only) TCP traffic, regardless of TCP port number.

f) Click **Add** to add the rule to the prefilter policy.
g) Click **Save** to save your changes to the prefilter policy.

You can now close the prefilter edit window and return to the access control policy edit window.

h) In the access control policy edit window, the Prefilter Policy dialog box should still be open. Click **OK** to save your changes to the prefilter policy assignment.
i) Click **Save** on the access control policy to save the changed prefilter policy assignment, if you changed it.

You can now deploy the changes to the affected devices.

Disable TCP Sequence Randomization

Each TCP connection has two initial sequence numbers (ISN): one generated by the client and one generated by the server. The Firewall Threat Defense device randomizes the ISN of the TCP SYN passing in both the inbound and outbound directions.

Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session. However, TCP sequence randomization effectively breaks TCP SACK (Selective Acknowledgement), as the sequence numbers the client sees are different from what the server sees.

You can disable TCP initial sequence number randomization if necessary, for example, because data is getting scrambled. Following are some situations where you might want to disable randomization.

- If another in-line firewall is also randomizing the initial sequence numbers, there is no need for both firewalls to be performing this action, even though this action does not affect the traffic.
- If you use eBGP multi-hop through the device, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum.
- If you use a WAAS device that requires the Firewall Threat Defense device not to randomize the sequence numbers of connections.
- If you enable hardware bypass for the ISA 3000, and TCP connections are dropped when the ISA 3000 is no longer part of the data path.

Procedure

Step 1 Create the extended ACL that defines the traffic class.

For example, to define a traffic class for TCP traffic from any host to 10.2.2.2, do the following:

a) Choose **Objects** > **Object Management**.

b) Choose **Access List** > **Extended** from the table of contents.
c) Click **Add Extended Access List**.
d) Enter a **Name** for the object, for example, preserve-sq-no.
e) Click **Add** to add a rule.
f) Keep **Allow** for the action.
g) Leave the **Source** list empty, enter 10.2.2.2 beneath the **Destination** list, and click **Add**.
h) Click **Port**, select **TCP (6)** beneath the **Selected Source Ports** list, and click **Add**. Do not enter a port number, simply add TCP as the protocol, which will cover all ports.
i) Click **Add** on the Extended Access List Entry dialog box to add the rule to the ACL.
j) Click **Save** on the Extended Access List Object dialog box to save the ACL object.

Step 2 Configure the service policy rule that disables TCP sequence number randomization.

For example, to disable randomization for this traffic class globally, do the following:

a) Choose **Policies** > **Access Control heading** > **Access Control**, and edit the policy assigned to the devices that require this service.
b) Click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line, and click **Edit** (✎) for the **Threat Defense Service Policy**.
c) Click **Add Rule**.
d) Select **Apply Globally** > **Next**.
e) Select the extended ACL object you created for this rule and click **Next**.
f) Deselect **Randomize TCP Sequence Number**.
g) (Optional.) Adjust the other connection options as needed.
h) Click **Finish** to add the rule. If necessary, drag and drop the rule to the desired position in the service policy.
i) Click **OK** to save the changes to the service policy.
j) Click **Save** on **Advanced** to save the changes to the access control policy.

You can now deploy the changes to the affected devices.

Examples for Service Policy Rules

The following topics provide examples of service policy rules.

Protect Servers from a SYN Flood DoS Attack (TCP Intercept)

A SYN-flooding denial of service (DoS) attack occurs when an attacker sends a series of SYN packets to a host. These packets usually originate from spoofed IP addresses. The constant flood of SYN packets keeps the server SYN queue full, which prevents it from servicing connection requests from legitimate users.

You can limit the number of embryonic connections to help prevent SYN flooding attacks. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination.

When the embryonic connection threshold of a connection is crossed, the Firewall Threat Defense acts as a proxy for the server and generates a SYN-ACK response to the client SYN request using the SYN cookie method, so that the connection is not added to the SYN queue of the targeted host. The SYN cookie is the initial sequence number returned in the SYN-ACK that is constructed from MSS, time stamp, and a mathematical hash of other items to essentially create a secret. If the Firewall Threat Defense receives an

ACK back from the client with the correct sequence number and within the valid time window, it can then authenticate that the client is real and allow the connection to the server. The component that performs the proxy is called TCP Intercept.

Setting connection limits can protect a server from a SYN flood attack. You can optionally enable TCP Intercept statistics and monitor the results of your policy. The following procedure explains the end-to-end process.

Before you begin

- Ensure that you set the embryonic connection limit lower than the TCP SYN backlog queue on the server that you want to protect. Otherwise, valid clients can no longer access the server during a SYN attack. To determine reasonable values for embryonic limits, carefully analyze the capacity of the server, the network, and server usage.
- Depending on the number of CPU cores on your Secure Firewall Threat Defense device model, the maximum concurrent and embryonic connections can exceed the configured numbers due to the way each core manages connections. In the worst case scenario, the device allows up to n-1 extra connections and embryonic connections, where n is the number of cores. For example, if your model has 4 cores, if you configure 6 concurrent connections and 4 embryonic connections, you could have an additional 3 of each type. To determine the number of cores for your model, enter the **show cpu core** command in the device CLI.

Procedure

Step 1 Create the extended ACL that defines the traffic class, which is the list of servers you want to protect.

For example, to define a traffic class to protect the web servers with the IP addresses 10.1.1.5 and 10.1.1.6:

a) Choose **Objects** > **Object Management**.
b) Choose **Access List** > **Extended** from the table of contents.
c) Click **Add Extended Access List**.
d) Enter a **Name** for the object, for example, protected-servers.
e) Click **Add** to add a rule.
f) Keep **Allow** for the action.
g) Leave the **Source** list empty, enter 10.1.1.5 beneath the **Destination** list, and click **Add**.
h) Also enter 10.1.1.6 beneath the **Destination** list and click **Add**.
i) Click **Port**, select **HTTP** in the available ports list, and click **Add to Destination**. If your server also support HTTPS connections, also add that port.
j) Click **Add** on the Extended Access List Entry dialog box to add the rule to the ACL.
k) Click **Save** on the Extended Access List Object dialog box to save the ACL object.

Step 2 Configure the service policy rule that sets embryonic connection limits.

For example, to set the total concurrent embryonic limit to 1000 connections, and the per-client limit to 50 connections, do the following:

a) Choose **Policies** > **Access Control heading** > **Access Control**, and edit the policy assigned to the devices that require this service.
b) Click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line, and click **Edit** (✎) for the **Threat Defense Service Policy**.

c) Click **Add Rule**.
d) Select **Apply Globally** > **Next**.
e) Select the extended ACL object you created for this rule and click **Next**.
f) Enter 1000 for **Connections** > **Maximum Embryonic**.
g) Enter 50 for **Connections Per Client** > **Maximum Embryonic**.
h) (Optional.) Adjust the other connection options as needed.
i) Click **Finish** to add the rule. If necessary, drag and drop the rule to the desired position in the service policy.
j) Click **OK** to save the changes to the service policy.
k) Click **Save** on **Advanced** to save the changes to the access control policy.

Step 3 (Optional.) Configure the rates for TCP Intercept statistics.

TCP Intercept uses the following options to determine the rate for collecting statistics. All options have default values, so if these rates suit your needs, you can skip this step.

- Rate Interval—The size of the history monitoring window, between 1 and 1440 minutes. The default is 30 minutes. During this interval, the system samples the number of attacks 30 times.
- Burst Rate—The threshold for syslog message generation, between 25 and 2147483647. The default is 400 per second. When the burst rate is exceeded, the device generates syslog message 733104.
- Average Rate—The average rate threshold for syslog message generation, between 25 and 2147483647. The default is 200 per second. When the average rate is exceeded, the device generates syslog message 733105.

If you want to adjust these options, do the following:

a) Choose **Objects** > **Object Management**.
b) Choose **FlexConfig** > **Text Object**.
c) Click **Edit** (✎) for the threat_defense_statistics system-defined object.
d) Although you can directly change the values, the recommended approach is to open the **Override** section and click **Add** to create a device override.
e) Select the devices to which you will assign the service policy (through the access control policy assignment) and click **Add** to move them to the selected list.
f) Click **Override**.
g) The object must have 3 entries, so click **Count** as needed until you get 3.
h) Enter the values you need, in order from 1-3, as rate interval, burst rate, and average rate. Consult the object description to verify you enter the values in the right order.
i) Click **Add** in the Object Override dialog box.
j) Click **Save** in the Edit Text Object dialog box.

Step 4 Enable TCP Intercept statistics.

You must configure a FlexConfig policy to enable TCP Intercept statistics.

a) Choose **Devices** > **FlexConfig**.
b) If you already have a policy assigned to the devices, edit it. Otherwise, create a new policy and assign it to the affected devices.
c) Select the **Threat_Detection_Configure** object in the **Available FlexConfig** list and click >>. The object is added to the **Selected Append FlexConfigs** list.
d) Click **Save**.

e) (Optional.) You can verify that you have the right settings by clicking **Preview Config** and selecting one of the devices.

The system generates the CLI commands that will be written to the device during the next deployment. These commands will include those needed for the service policy as well as those needed for threat detection statistics. Scroll to the bottom of the preview to see the appended CLI. The TCP Intercept statistics command should look something like the following, if you use the default values (line break added for clarity):

```
###Flex-config Appended CLI ###

threat-detection statistics tcp-intercept rate-interval 30
       burst-rate 400 average-rate 200
```

Step 5 You can now deploy the changes to the affected devices.

Step 6 Monitor the TCP Intercept statistics from the device CLI using the following commands:

- **show threat-detection statistics top tcp-intercept** [**all** | **detail**]—View the top 10 protected servers under attack. The **all** keyword shows the history data of all the traced servers. The **detail** keyword shows history sampling data. The system samples the number of attacks 30 times during the rate interval, so for the default 30 minute period, statistics are collected every 60 seconds.

  Note
  You can use the **shun** command to block attacking host IP addresses. To remove the block, use the **no shun** command.

- **clear threat-detection statistics tcp-intercept**—Erases TCP Intercept statistics.

Example:

```
hostname(config)# show threat-detection statistics top tcp-intercept
Top 10 protected servers under attack (sorted by average rate)
Monitoring window size: 30 mins    Sampling interval: 30 secs
<Rank> <Server IP:Port> <Interface> <Ave Rate> <Cur Rate> <Total> <Source IP (Last Attack
Time)>
--------------------------------------------------------------------------------
1    10.1.1.5:80 inside 1249 9503 2249245 <various> Last: 10.0.0.3 (0 secs ago)
2    10.1.1.6:80 inside 10 10 6080 10.0.0.200 (0 secs ago)
```

Make the Firewall Threat Defense Device Appear on Traceroutes

By default, the Firewall Threat Defense device does not appear on traceroutes as a hop. To make it appear, you need to decrement the time-to-live on packets that pass through the device, and increase the rate limit on ICMP unreachable messages. To accomplish this, you must configure a service policy rule and adjust the ICMP platform settings policy.

Note If you decrement time to live, packets with a TTL of 1 will be dropped, but a connection will be opened for the session on the assumption that the connection might contain packets with a greater TTL. Note that some packets, such as OSPF hello packets, are sent with TTL = 1, so decrementing time to live can have unexpected consequences. Keep these considerations in mind when defining your traffic class.

Procedure

Step 1 Create the extended ACL that defines the traffic class for which to enable traceroute reporting.

For example, to define a traffic class for all addresses, but excluding OSPF traffic, do the following:

a) Choose **Objects** > **Object Management**.
b) Choose **Access List** > **Extended** from the table of contents.
c) Click **Add Extended Access List**.
d) Enter a **Name** for the object, for example, traceroute-enabled.
e) Click **Add** to add a rule to exclude OSPF.
f) Change the action to **Block**, click **Port**, select **OSPFIGP (89)** as the protocol beneath the **Destination Ports** list, and click **Add** to add the protocol to the selected list.
g) Click **Add** on the Extended Access List Entry dialog box to add the OSPF rule to the ACL.
h) Click **Add** to add a rule to include all other connections.
i) Keep **Allow** for the action, and leave both the Source and Destination lists empty.
j) Click **Add** on the Extended Access List Entry dialog box to add the rule to the ACL.

   Ensure that the OSPF deny rule is above the Allow Any rule. Drag and drop to move the rules if necessary.

k) Click **Save** on the Extended Access List Object dialog box to save the ACL object.

Step 2 Configure the service policy rule that decrements the time-to-live value.

For example, to decrement time-to-live globally, do the following:

a) Choose **Policies** > **Access Control heading** > **Access Control**, and edit the policy assigned to the devices that require this service.
b) Click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line, and click **Edit** (✎) for the **Threat Defense Service Policy**.
c) Click **Add Rule**.
d) Select **Apply Globally** and click **Next**.
e) Select the extended ACL object you created for this rule and click **Next**.
f) Select **Enable Decrement TTL**.
g) (Optional.) Adjust the other connection options as needed.
h) Click **Finish** to add the rule. If necessary, drag and drop the rule to the desired position in the service policy.
i) Click **OK** to save the changes to the service policy.
j) Click **Save** on **Advanced** to save the changes to the access control policy.

   You can now deploy the changes to the affected devices.

Step 3 Increase the rate limit on ICMP unreachable messages.

a) Choose **Devices** > **Platform Settings**.
b) If you already have a policy assigned to the devices, edit it. Otherwise, create a new Threat Defense platform settings policy and assign it to the affected devices.
c) Select **ICMP** from the table of contents.
d) Increase the **Rate Limit**, for example, to 50. You might also want to increase the **Burst Size**, for example, to 10, to ensure enough responses are generated within the rate limit.

   You can leave the ICMP rules table empty, it is not related to this task.

e) Click **Save**.

Step 4 You can now deploy the changes to the affected devices.

Monitoring Service Policies

You can monitor service-policy related information using the device CLI. Following are some useful commands.

- **show conn** [**detail**]

  Shows connection information. Detailed information uses flags to indicate special connection characteristics. For example, the "b" flag indicates traffic subject to TCP State Bypass.

  When you use the **detail** keyword, you can see information about Dead Connection Detection (DCD) probing, which shows how often the connection was probed by the initiator and responder. For example, the connection details for a DCD-enabled connection would look like the following:

```
TCP dmz: 10.5.4.11/5555 inside: 10.5.4.10/40299,
    flags UO , idle 1s, uptime 32m10s, timeout 1m0s, bytes 11828,
cluster sent/rcvd bytes 0/0, owners (0,255)
  Traffic received at interface dmz
        Locally received: 0 (0 byte/s)
  Traffic received at interface inside
        Locally received: 11828 (6 byte/s)
  Initiator: 10.5.4.10, Responder: 10.5.4.11
  DCD probes sent: Initiator 5, Responder 5
```

- **show service-policy**

  Shows service policy statistics, including Dead Connection Detection (DCD) statistics.

- **show threat-detection statistics top tcp-intercept** [**all** | **detail**]

  View the top 10 protected servers under attack. The **all** keyword shows the history data of all the traced servers. The **detail** keyword shows history sampling data. The system samples the number of attacks 30 times during the rate interval, so for the default 30 minute period, statistics are collected every 60 seconds.

History for Threat Defense Service Policy

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Description
Threat Defense Service Policy.	6.3	Any	You can now configure a Threat Defense Service Policy as part of your access control policy advanced options. You can use Threat Defense Service Policies to apply services to specific traffic classes. Features supported include TCP State Bypass, randomizing TCP sequence numbers, decrementing the time-to-live (TTL) value on packets, Dead Connection Detection, setting a limit on the maximum number of connections and embryonic connections per traffic class and per client, and timeouts for embryonic, half closed, and idle connections. New screen: **Policies** > **Access Control** > **Access Control**, **Advanced** tab, **Threat Defense Service Policy**. Supported platforms: Secure Firewall Threat Defense
Initiator and responder information for Dead Connection Detection (DCD), and DCD support in a cluster.	6.5	Any	If you enable Dead Connection Detection (DCD), you can use the **show conn detail** command to get information about the initiator and responder. Dead Connection Detection allows you to maintain an inactive connection, and the **show conn** output tells you how often the endpoints have been probed. In addition, DCD is now supported in a cluster. New/Modified commands: **show conn** (output only). Supported platforms: Secure Firewall Threat Defense
Configure the maximum segment size (MSS) for embryonic connections.	7.1	Any	You can configure a service policy to set the server maximum segment size (MSS) for SYN-cookie generation for embryonic connections upon reaching the embryonic connections limit. This is meaningful for service policies where you are also setting embryonic connection maximums. New or changed screens: **Connection Settings** in the Add/Edit Service Policy wizard.

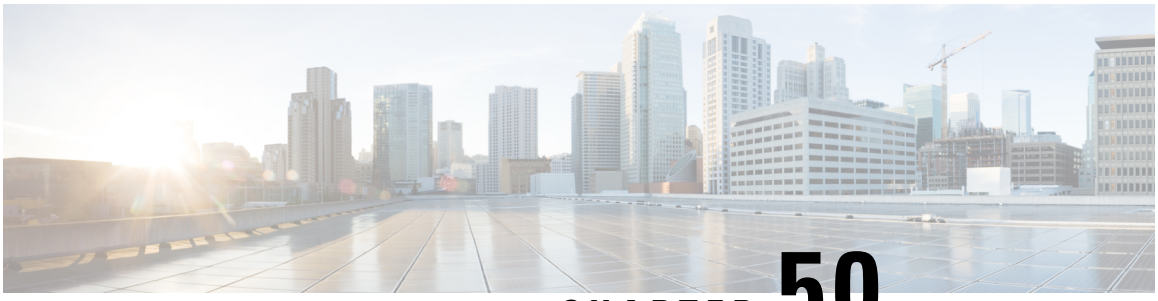

CHAPTER 50

Threat Detection

Threat Detection's portscan detector is a mechanism designed to help you detect and prevent portscan activity in all types of traffic to protect networks from eventual attacks. Portscan traffic can be detected efficiently in both allowed and denied traffic.

Portscan is a form of network reconnaissance that is often used by attackers as a prelude to an attack. In a portscan, an attacker determines the types of network protocols or services a host supports and sends specially crafted packets to a targeted host. By examining the packets that the host responds with, the attacker can often determine which ports are open on the host and, either directly or by inference, which application protocols are running on these ports.

- Portscan Detection and Prevention, on page 1955
- Best Practices for Portscan Prevention, on page 1957
- Requirements and Prerequisites for Threat Detection, on page 1958
- Guidelines and Limitations for Threat Detection, on page 1958
- Configure Portscan Detection and Prevention, on page 1959
- Monitoring Threat Detection, on page 1961
- History for Threat Detection, on page 1963

Portscan Detection and Prevention

Use Threat Detection to identify port scan activity. You can use the system to detect port scans and issue events when they are found. Optionally, you can configure the system to prevent port scans by automatically blocking scanners. When preventing port scans, the system sends you events and also blocks the attacker for a duration period that you set.

Pre-Defined Sensitivity Levels for Portscan Detection

When configuring detection settings, you select from the following pre-defined sensitivity levels. Except for Custom, each level has pre-set values for each protocol for the number of ports (TCP/UDP), protocols (IP), or hosts (TCP/UDP/IP/ICMP) that must be scanned within a set time interval (in seconds). Also, all types of scan/sweep are enabled.

Note When counting ports/protocols, Threat Detection increments the number if the port/protocol in the current packet differs from the previous packet. For example, if you have an application that opens connections in 10 set ports randomly, the total number of ports scanned could mount so quickly that your port number will be exceeded within the interval. The system does not count unique ports only.

Exceeding the number within the interval can indicate a scanning attack. Portscan events are generated only when the port/protocol/host numbers are exceeded for the moving time interval window.

- **Low**—This level uses the shortest time window for portscan detection, coupled with high counts for port/protocol/host. Thus, you should see portscan events for the most aggressive scanners only. Select this sensitivity level to suppress false positives, but remember that some types of portscans (slow scans, filtered scans) might be missed. For more detail on how low sensitivity detection works, see Detection in the Low Sensitivity Level, on page 1957.
  - **Interval** (TCP/UDP/IP/ICMP)—60 seconds.
  - **TCP/UDP portscan Number of Ports**—120.
  - **TCP/UDP portsweep Number of Hosts**—180.
  - **IP protocol scan Number of Protocols**—30.
  - **IP protocol sweep Number of Hosts**—25.
  - **ICMP host sweep Number of Hosts**—50.
- **Medium**—This level uses moderate values for both the interval and port/protocol/host counts. However, very active hosts such as network address translators and proxies might generate false positives. Add such hosts to the ignore scanner list. This is the default sensitvity level and a good place to start.
  - **Interval** (TCP/UDP/IP/ICMP)—90 seconds.
  - **TCP/UDP portscan Number of Ports**—90.
  - **TCP/UDP portsweep Number of Hosts**—150.
  - **IP protocol scan Number of Protocols**—15.
  - **IP protocol sweep Number of Hosts**—20.
  - **ICMP host sweep Number of Hosts**—30.
- **High**—This level uses a much longer time window for portscan detection, coupled with lower counts for port/protocol/host. With this level, you are most likely to see events for even the least aggressive port scans/sweeps, so you are more likely to notice all attackers. On the other hand, this level would likely result in the most portscan events issued, and potentially the highest number of false positives.
  - **Interval** (TCP/UDP/IP/ICMP)—600 seconds (10 minutes).
  - **TCP/UDP portscan Number of Ports**—60.
  - **TCP/UDP portsweep Number of Hosts**—100.
  - **IP protocol scan Number of Protocols**—10.
  - **IP protocol sweep Number of Hosts**—10.

- **ICMP host sweep Number of Hosts**—20.

- **Custom**—If you want to configure any setting differently than one of the pre-defined sensitivity levels, or disable a particular type of scan/sweep, the level automatically switches to custom. If you want to adjust the options, first select the level that most closely matches what you want, then edit the values as appropriate.

Detection in the Low Sensitivity Level

If you select the low sensitivity level, the system tracks negative responses for TCP, UDP, and ICMP initial packets. Only if the number of unsuccessful connections is more than the rejection threshold (10% in low sensitivity) and the port/IP protocol count is more than the configured threshold, is an alert triggered. This mitigates false positives.

Rejection threshold applies to low sensitivity (or equivalent custom settings) only; it does not apply to other sensitivity levels or their custom equivalents.

If there is a mix of allowed and blocked traffic, the number of rejected ports or hosts is calculated based on the difference between allowed and blocked traffic. In the case of only blocked traffic, the rejection threshold is not considered.

These criteria are not used for UDP/ICMP connections on interfaces configured in inline sets.

For example, in low sensitivity mode, the port count threshold is 120. Thus, the rejection count threshold is 10% of 120, which is 12. Following are examples of how the system would issues portscan events under this configuration:

- An attacker initiates connections with 131 ports of the target and the target positively acknowledges all the initiations. Port count = 131, which is greater than the threshold, but a portscan alert is not triggered because there are no negative acknowledgements.
- An attacker initiates connections with 131 ports of the target and the target positively acknowledges 121 initiations and negatively acknowledges 10 initiations. Port count = 131, which is greater than the threshold, but reject port count = 10, which is lesser than the rejection threshold. Therefore, a portscan alert is not triggered.
- An attacker initiates connections with 134 ports of the target and the target positively acknowledges 121 initiations and negatively acknowledges 13 initiations. Port count = 134, which is greater than the threshold, reject port count = 13 is also higher than the rejection threshold. Therefore a portscan alert is triggered.

Best Practices for Portscan Prevention

Portscan prevention mode can result in unintended traffic outage. In Prevention mode, hosts are blocked from further scanning of networks on all protocols for the configured duration. Review the detection and prevention parameters carefully to ensure legitimate traffic is not blocked.

Before configuring portscan in Prevention mode, we strongly recommend the following:

1. Start using portscan in Detection mode.
2. Observe the generated portscan events.

3. Tune the Sensitivity level, and Monitored networks, Ignore Scanner list, and Ignore Target list. If a pre-defined sensitivity level does not work well for your situation, configure custom settings as needed.
4. Repeat the process until false positives are eliminated and the event rate represents an accurate picture of port scanning in your network. Ensure that you are comfortable with blocking the remaining identified scanners.

Requirements and Prerequisites for Threat Detection

Model Support

Threat Defense running version 7.2+ and Snort 3.

Supported Domains

Any

User Roles

Admin

Access Admin

Network Admin

Guidelines and Limitations for Threat Detection

- Threat detection requires Snort 3. The managed device must be at version 7.2 or higher. For Snort 2, or devices at versions lower than 7.2, you can configure port scan through the NAP policy. Note that the Threat Detection feature is not the same as the port scan feature in the NAP policy. If there are non-Snort 3/version 7.2+ devices assigned to the access control policy, the Threat Detection settings will not be deployed to those unsupported devices.
- If you configure port scan in the NAP policy on a device running 7.1 or lower, that configuration is not translated to the Threat Detection feature on upgrade to 7.2. You must manually configure Threat Detection. Although the NAP and Threat Detection portscan options are similar, they do not match one-to-one.
- If you configure Threat Detection, any port scan configuration in the NAP policy is ignored and not configured on the devices that support Threat Detection.
- The NAP port scan feature for Snort 3 is always ignored for version 7.2+ devices. To configure port scanning, you must use the Threat Defense settings.
- Threat detection works on traffic that passes through the device only. It does not work on traffic directed to the device.
- In a high availability setup, port scanning statistics are not synchronized to the standby unit. However, blocked hosts are synchronized and continue to be blocked until the duration period expires in case of a failover.

- (Devices running Threat Defense versions 7.2-7.7). For nodes in a cluster, detection and prevention happen on the individual cluster node. That is, if node B detects and blocks traffic from a host, node A will not be aware of that action because port scan statistics are not synchronized across cluster nodes.
- For inline sets, or for interfaces that are configured to be part of an Equal-Cost Multipath (ECMP) traffic zone, detection and prevention are done at the zone level. Portscan statistics for a host are accumulated across all interfaces of a zone. Similarly, when a host crosses configured thresholds, it is blocked across all interfaces of the corresponding zone.
- Although the portscan events generated by the Threat Detection feature are the same as the ones Snort issues for port scan, you do not need to enable port scanning intrusion rules to get the events. Threat Detection works regardless of your intrusion policy implementation.

Configure Portscan Detection and Prevention

Portscan is a form of network reconnaissance that is often used by attackers as a prelude to an attack. In a portscan, an attacker determines the types of network protocols or services a host supports and sends specially crafted packets to a targeted host. By examining the packets that the host responds with, the attacker can often determine which ports are open on the host and, either directly or by inference, which application protocols are running on these ports.

You can enable Threat Detection to watch for port scanning activity and optionally, automatically block scanners for a period of time.

Before you begin

FQDN, Wildcard mask, any, any-ipv4, and any-ipv6 network objects are not supported for portscan configuration. These objects are not shown in the **Monitor**, **Ignore Scanner**, **Ignore Target**, and **Exclude** fields.

Procedure

Step 1 In the access control policy editor, click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line. Then, click **Edit** (✎) next to **Threat Detection**.

Step 2 In the **Threat Detection** window, select the **Portscan mode**:

- **Disable**—Turn off Threat Detection. This is the default mode. You can click **Revert to Defaults** to return to this unconfigured state.
- **Detection**—Perform portscan detection, but alert on problems only. Do not take action against potential attackers. We suggest you use this mode initially until you fine-tune the Threat Detection settings to avoid excessive false positives.
- **Prevention**—Perform portscan detection and actively block identified scanners, that is, hosts that are performing the port scan.

Step 3 Configure the **Traffic Selection** options.

The traffic selection options determine which networks are monitored, the type of connections monitored, and whether any scanners or target hosts should be exempted from the monitored networks. By default, the system monitors permitted connections on all networks.

- **Detection On Traffic**—Select the types of connection that will be monitored for portscan activity: **Permitted**, **Denied**, or **All** traffic. The default is **Permitted**.
- **Monitor**—Select the network objects that define the networks to monitor for portscan or sweep activity. The default is any network, IPv4 or IPv6. Use this option to limit scanning to untrusted networks.
- **Ignore Scanner**—Select the network objects that define the hosts or networks, from within the range of the monitored networks, that should be ignored. For example, if you have set up your own scanner to test your network, you can exempt the address of your scanner to avoid unnecessary reporting on the address. Do not include addresses that are outside the monitored networks, as these addresses are already ignored.
- **Ignore Target**—Select the network objects that define the hosts or networks that should be ignored as targets, that is, victims of a portscan or sweep.

Step 4 Click the **Configuration** tab and select the scanning sensitivity level.

The pre-defined sensitivity levels, **Low**, **Medium**, and **High**, set the port scanning options to values that are increasingly aggressive. For example, if you select Low, you would expect to see fewer port scanning events, and you could potentially miss attackers more easily than if you selected Medium or High. On the other hand, if you select High, you might see more events and also potentially more false positives. The default level is Medium. For more information on these levels, see Pre-Defined Sensitivity Levels for Portscan Detection, on page 1955.

As you select the levels, you can see the related values within the protocol sections: **TCP**, **UDP**, **IP**, and **ICMP**. If you change any of the preset values, or disable a type of scan, the sensitivity mode automatically changes to **Custom**.

Within each protocol section, the options are:

- **Interval**—The time range, in seconds, within which the configured values for portscan or portsweep are exceeded. For example, if you select 90 seconds, and 60 as the number of TCP portscan ports, a scanner would need to try 60 ports on a host within 90 seconds for it to be considered a portscan. The system generates events only if the number of ports, protocols, or hosts (for a portsweep) are exceeded within the specified interval.

  You can specify a range between 30-600 seconds. The longer the period, the more likely a host might be identified as a scanner.
- **Portscan** (TCP/UDP)—Select whether to monitor for port scanning against single hosts, and specify the number of ports that must be scanned within the interval to count as a portscan attack. The allowed range is 1-256.
- **Portsweep** (TCP/UDP)—Select whether to monitor for port sweeping against multiple hosts, and specify the number of hosts that must be scanned for a given port within the interval to count as a portsweep attack. The allowed range is 1-256.
- **Protocol Scan** (IP)—Select whether to monitor for protocol scanning against single hosts, and specify the number of protocols that must be scanned within the interval to count as a protocol scan attack. The allowed range is 1-255.

- **Protocol Sweep** (IP)—Select whether to monitor for protocol sweeping against multiple hosts, and specify the number of hosts that must be scanned for a given protocol within the interval to count as a protocol sweep attack. The allowed range is 1-256.
- **Hostsweep** (ICMP)—Select whether to monitor for ICMP host sweeping against multiple hosts, and specify the number of hosts that must be scanned within the interval to count as a hostsweep attack. The allowed range is 1-256.

Step 5 If you selected prevention mode, click the **Prevention** tab and configure the options.

In Prevention mode, hosts are automatically blocked from further scanning of networks on all protocols for the configured duration. Review the detection and prevention parameters carefully to ensure legitimate traffic is not blocked.

- **Exclude**—Select the network objects that define the hosts or networks, from within the range of the monitored networks, that should be excluded from automatic blocking. Even if these hosts voilate your scanning detection parameters, the system will not block them.
- **Duration**—How long, in seconds, automatically blocked scanner hosts should be prevented from sending traffic of any kind through the device. After the duration period ends, the hosts are automatically cleared and can again send traffic through the device. The allowed range is 600-2592000 seconds. The default is 3600 seconds (1 hour).

  If you need to manually unblock a host, SSH to the firewall that is blocking the host and use the **clear threat-detection portscan attacker** command.

Step 6 Click **OK** to save the Threat Detection settings.

Step 7 Click **Save** to save the access control policy.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Monitoring Threat Detection

The following topics explain how to monitor portscan activity

Viewing Portscan Alerts

Portscan activity is alerted through the existing portscan-specific intrusion events. Intrusion events with Generator ID (GID) 122 and Snort ID from SIDs 1 through 27 are generated. For these events, the *(port_scan)* string is prepended in the event messages. The events include packet information along with packet data containing the statistics that triggered the alert.

To see portscan events, go to **Analysis** > **Intrusion** > **Events**.

Portscan issues these events regardless of your instrusion policy or NAP configuration. Events are issued only when scanners exceed the number of configured ports/protocols/hosts for the various types of scan or sweep within the configured time interval for the associated protocol. A port scan from one host generates one event

per set interval as soon as the threshold is met. If the same host initiates a new port scan during the same interval, no event is reported.

The following table shows the possible events.

Table 106: Portscan Events

Portscan Type	Intrusion Event
TCP Regular, Decoy, Distributed Scan	122:1 (port_scan) TCP portscan
TCP Portsweep	122:3 (port_scan) TCP portsweep
TCP Distributed Scan	122:4 (port_scan) TCP distributed portscan
IP Regular, Decoy, Distributed Protocol Scan n	122:9 (port_scan) IP protocol scan
IP Protocol Sweep	122:11 (port_scan) IP protocol sweep
IP Distributed Scan	122:12 (port_scan) IP distributed protocol scan
UDP Regular, Decoy, Distributed Scan	122:17 (port_scan) UDP portscan
UDP Portsweep	122:19 (port_scan) UDP portsweep
UDP Distributed Scan	122:20 (port_scan) UDP distributed portscan
ICMP Sweep	122:25 (port_scan) ICMP sweep
Portscan Block	122:100 (port_scan) host blocked due to portscan activity

Monitoring Portscan on the Firewall

To monitor portscan, log into the device CLI and use the following commands.

- **show threat-detection portscan** [**attacker** | **target** | **shun**]

  Shows the IP addresses of scanners, those that have been shunned (blocked), and hosts that have been targeted by scans or sweeps.

- **show threat-detection portscan statistics** [**host** [*ipv4_address* | *ipv6_address*]] [**protocol** {**tcp** | **udp** | **ip** | **icmp**}]

  Shows statistics related to the portscan system. You can specify host, protocol, or host and protocol to filter the output to the desired information.

- **clear threat-detection portscan** [**attacker** | **target** | **shun**] [*ipv4_address mask* | *ipv6_address/prefix*]

  Manually unblocks scanners (attackers) or identified targets. Enter the command without parameters to clear all attackers, targets, or shunned hosts.

- **clear threat-detection portscan statistics** [**host** [*ipv4_address* | *ipv6_address*]] [**protocol** {**tcp** | **udp** | **ip** | **icmp**}]

Erases statistics related to portscan, so that you can more clearly see the current state of scanning through this device. Enter the command without parameters to clear all statistics. Alternatively, specify a host, protocol, or host and protocol, to limit the reset to the specified items.

Unblocking A Host

If you configure Threat Detection in prevention mode, and the system blocks a host that you know is not an attacker, you can manually unblock the host before host is automatically unblocked when the duration period expires.

To manually unblock the host, log into the device CLI where the host is blocked and enter the **clear threat-detection portscan attacker** command. For example:

```
> clear threat-detection portscan attacker 10.2.0.100 255.255.255.255
1 tracker object deleted and 1 shun entry removed
```

Consider adding the host IP to the Exclude list in the prevention configuration.

History for Threat Detection

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Description
Improved detection in low sensitivity.	7.4	7.4	The method for identifying port scans and sweeps has been improved when operating in low sensitivity level. Changes are automatic; there are no new configuration settings.
Improved portscan detection.	7.2	7.2 running Snort 3	With an improved portscan detector, you can easily configure the system to detect or prevent portscans. You can refine the networks you want to protect, set the sensitivity, and so on. For devices running Snort 2 and for devices running Version 7.1 and earlier, continue to use the network analysis policy for portscan detection. New/modified screens: We added **Threat Detection** to the access control policy's Advanced tab. New/Modified commands: **clear threat-detection portscan**, **show threat-detection portscan**.

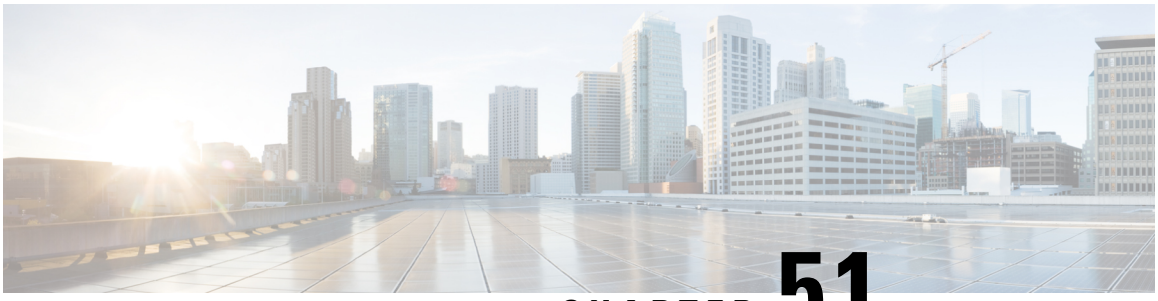

CHAPTER 51

Intelligent Application Bypass

The following topics describe how to configure access control polices to use Intelligent Application Bypass (IAB)

- Introduction to IAB, on page 1965
- IAB Options, on page 1966
- Requirements and Prerequisites for Intelligent Application Bypass, on page 1968
- Configuring Intelligent Application Bypass, on page 1968
- IAB Logging and Analysis, on page 1969

Introduction to IAB

IAB identifies applications that you trust to traverse your network without further inspection if performance and flow thresholds are exceeded. For example, if a nightly backup significantly impacts system performance, you can configure thresholds that, if exceeded, trust traffic generated by your backup application. Optionally, you can configure IAB so that, when an inspection performance threshold is exceeded, IAB trusts all traffic that exceeds any flow bypass threshold, regardless of the application type.

The system implements IAB on traffic allowed by access control rules or the access control policy's default action, before the traffic is subject to deep inspection. A test mode allows you to determine whether thresholds are exceeded and, if so, to identify the application flows that would have been bypassed if you had actually enabled IAB (called *bypass mode*).

The following graphic illustrates the IAB decision-making process:

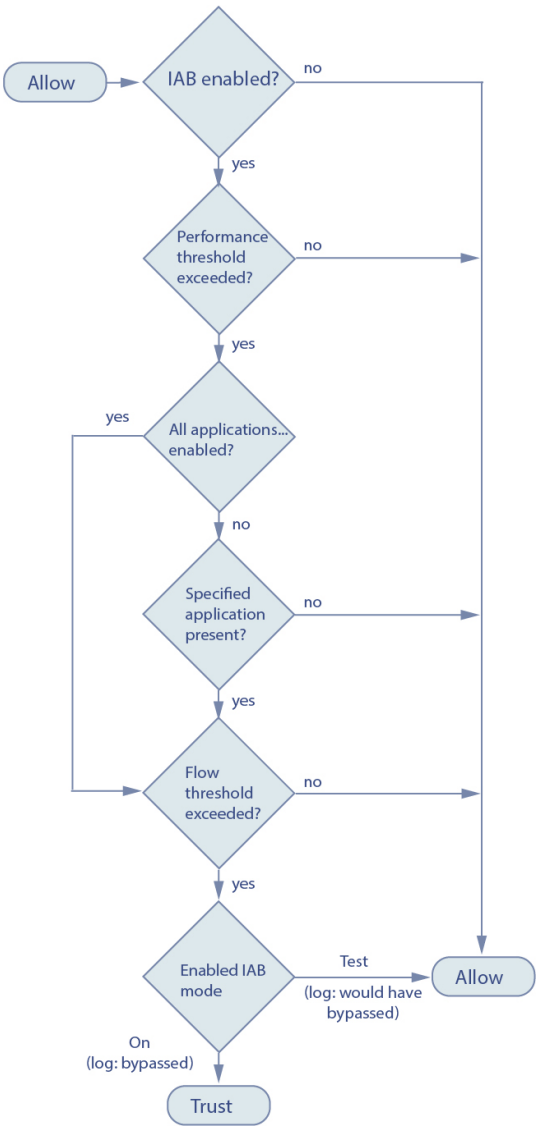

IAB Options

State

Enables or disables IAB.

Performance Sample Interval

Specifies the time in seconds between IAB performance sampling scans, during which the system collects system performance metrics for comparison to IAB performance thresholds. A value of `0` disables IAB.

Bypassable Applications and Filters

This feature provides two mutually exclusive options:

Applications/Filters
Provides an editor where you can specify bypassable applications and sets of applications (filters). See Application Rule Conditions, on page 939.

All applications including unidentified applications
When an inspection performance threshold is exceeded, trusts all traffic that exceeds any flow bypass threshold, regardless of the application type.

Performance and Flow Thresholds

You must configure at least one inspection performance threshold and one flow bypass threshold. When a performance threshold is exceeded, the system examines flow thresholds and, if one threshold is exceeded, trusts the specified traffic. If you enable more than one of either, only one of each must be exceeded.

Inspection performance thresholds provide intrusion inspection performance limits that, if exceeded, trigger the inspection of flow thresholds. IAB does not use inspection performance thresholds set to `0`. You can configure one or more of the following inspection performance thresholds:

Drop Percentage

Average packets dropped as a percentage of total packets, when packets are dropped because of performance overloads caused by expensive intrusion rules, file policies, decompression, and so on. This does not refer to packets dropped by normal configurations such as intrusion rules. Note that specifying an integer greater than 1 activates IAB when the specified percentage of packets is dropped. When you specify `1`, any percentage from 0 through 1 activates IAB. This allows a small number of packets to activate IAB.

Processor Utilization Percentage

Average percentage of processor resources used.

Packet Latency

Average packet latency in microseconds.

Flow Rate

The rate at which the system processes flows, measured as the number of flows per second. Note that this option configures IAB to measure flow *rate*, not flow *count*.

Flow bypass thresholds provide flow limits that, if exceeded, trigger IAB to trust bypassable application traffic in bypass mode or allow application traffic subject to further inspection in test mode. IAB does not use flow bypass thresholds set to `0`. You can configure one or more of the following flow bypass thresholds:

Bytes per Flow

The maximum number of kilobytes a flow can include.

Packets per Flow

The maximum number of packets a flow can include.

Flow Duration

The maximum number of seconds a flow can remain open.

Flow Velocity

The maximum transfer rate in kilobytes per second.

Requirements and Prerequisites for Intelligent Application Bypass

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Configuring Intelligent Application Bypass

Caution Not all deployments require IAB, and those that do might use it in a limited fashion. Do not enable IAB unless you have expert knowledge of your network traffic, especially application traffic, and system performance, including the causes of predictable performance issues. Before you run IAB in bypass mode, make sure that trusting the specified traffic does not expose you to risk.

Before you begin

For Classic devices, you must have the Control license.

Procedure

Step 1 In the access control policy editor, click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line. Then, click **Edit** ( ) next to **Intelligent Application Bypass Settings**.

Step 2 Configure IAB options:

- **State**—Turn IAB **Off** or **On**, or enable IAB in **Test** mode.
- **Performance Sample Interval**—Enter the time in seconds between IAB performance-sampling scans. If you enable IAB, even in test mode, enter a non-zero value. Entering `0` disables IAB.
- **Bypassable Applications and Filters**—Choose from:
  - Click the number of bypassed applications and filters and specify the applications whose traffic you want to bypass; see Configuring Application Conditions and Filters, on page 1789.

- Click **All applications including unidentified applications** so that, when an inspection performance threshold is exceeded, IAB trusts all traffic that exceeds any flow bypass threshold, regardless of the application type.

- **Inspection Performance Thresholds**—Click **Configure** and enter at least one threshold value.
- **Flow Bypass Thresholds**—Click **Configure** and enter at least one threshold value.

You must specify at least one inspection performance threshold and one flow bypass threshold; both must be exceeded for IAB to trust traffic. If you enter more than one threshold of each type, only one of each type must be exceeded. For detailed information, see IAB Options, on page 1966.

Step 3 Click **OK** to save IAB settings.

Step 4 Click **Save** to save the policy.

What to do next

- Because some packets must be allowed to pass before an application can be detected, you must configure your system to examine those packets.
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

IAB Logging and Analysis

IAB forces an end-of-connection event that logs bypassed flows and flows that would have been bypassed, regardless of whether you have enabled connection logging. Connection events indicate flows that are bypassed in bypass mode or that would have been bypassed in test mode. Custom dashboard widgets and reports based on connection events can display long-term statistics for bypassed and would-have-bypassed flows.

IAB Connection Events

Action

When **Reason** includes `Intelligent App Bypass`:

`Allow` -

indicates that the applied IAB configuration was in test mode and traffic for the application specified by **Application Protocol** remains available for inspection.

`Trust` -

indicates that the applied IAB configuration was in bypass mode and traffic for the application specified by **Application Protocol** has been trusted to traverse the network without further inspection.

Reason

`Intelligent App Bypass` indicates that IAB triggered the event in bypass or test mode.

Application Protocol

This field displays the application protocol that triggered the event.

Example

In the following truncated graphic, some fields are omitted. The graphic shows the **Action**, **Reason**, and **Application Protocol** fields for two connection events resulting from different IAB settings in two separate access control policies.

For the first event, the `Trust` action indicates that IAB was enabled in bypass mode and Bonjour protocol traffic was trusted to pass without further inspection.

For the second event, the `Allow` action indicates that IAB was enabled in test mode, so Ubuntu Update Manager traffic was subject to further inspection but would have been bypassed if IAB had been in bypass mode.

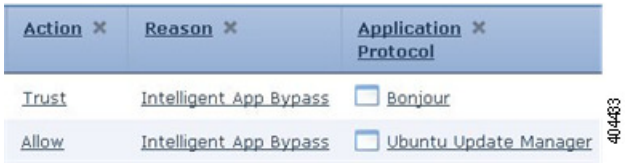

Example

In the following truncated graphic, some fields are omitted. The flow in the second event was both bypassed (**Action**: `Trust`; **Reason**: `Intelligent App Bypass`) and inspected by an intrusion rule (**Reason**: `Intrusion Monitor`). The `Intrusion Monitor` reason indicates that an intrusion rule set to **Generate Events** detected but did not block an exploit during the connection. In the example, this happened before the application was detected. After the application was detected, IAB recognized the application as bypassable and trusted the flow.

IAB Custom Dashboard Widgets

You can create a Custom Analysis dashboard widget to display long-term IAB statistics based on connection events. Specify the following when creating the widget:

- **Preset**: `None`
- **Table**: `Application Statistics`
- **Field**: any
- **Aggregate**: either of:
  - `IAB Bypassed Connections`
  - `IAB Would Bypass Connections`
- **Filter**: any

Examples

In the following Custom Analysis dashboard widget examples:

- The *Bypassed* example shows statistics for application traffic bypassed because the applications were specified as bypassable and IAB was enabled in bypass mode in the deployed access control policy.
- The *Would Have Bypassed* example shows statistics for application traffic that would have been bypassed because the applications were specified as bypassable and IAB was enabled in test mode in the deployed access control policy. .

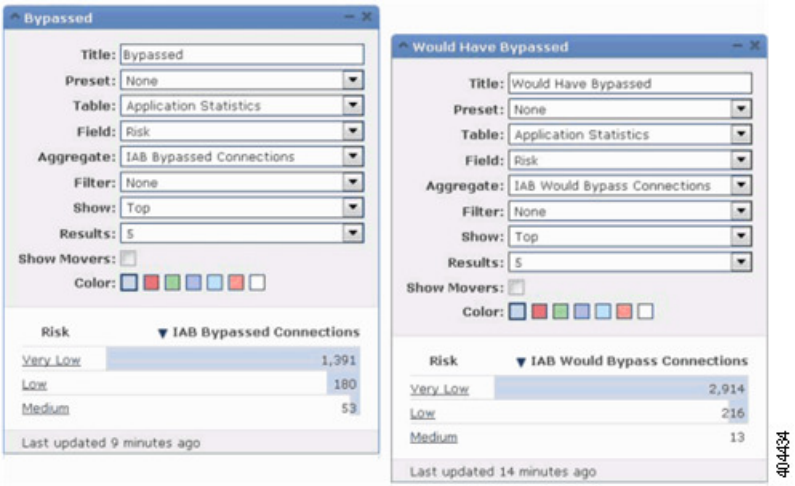

IAB Custom Reports

You can create a custom report to display long-term IAB statistics based on connection events. Specify the following when creating the report:

- **Table**: `Application Statistics`
- **Preset**: `None`
- **Filter**: any
- **X-Axis**: any
- **Y-Axis**: either of:
  - `IAB Bypassed Connections`
  - `IAB Would Bypass Connections`

Examples

The following graphic shows two abbreviated report examples:

- The *Bypassed* example shows statistics for application traffic bypassed because the applications were specified as bypassable and IAB was enabled in bypass mode in the deployed access control policy.
- The *Would Have Bypassed* example shows statistics for application traffic that would have been bypassed because the applications were specified as bypassable and IAB was enabled in test mode in the deployed access control policy.

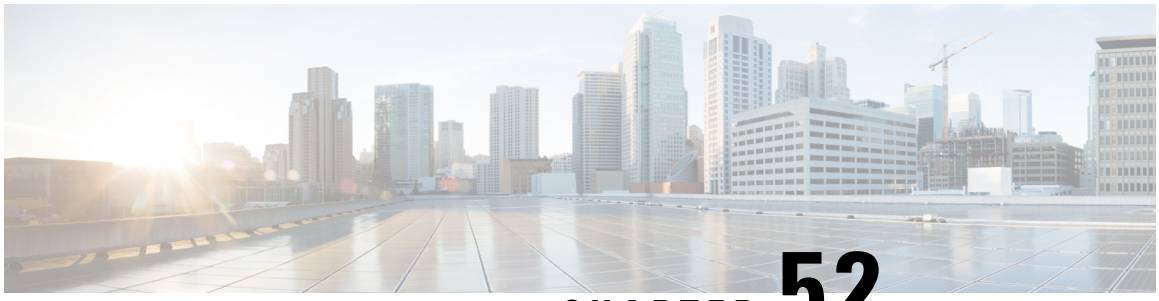

CHAPTER 52

Content Restriction

The following topics describe how to configure access control policies to use content restriction features:

- About Content Restriction, on page 1973
- Requirements and Prerequisites for Content Restriction, on page 1974
- Guidelines and Limitations for Content Restriction, on page 1975
- Using Access Control Rules to Enforce Content Restriction, on page 1975
- Using a DNS Sinkhole to Enforce Content Restriction, on page 1976

About Content Restriction

Major search engines and content delivery services provide features that allow you to restrict search results and website content. For example, schools use content restriction features to comply with the Children's Internet Protection Act (CIPA).

When implemented by search engines and content delivery services, you can enforce content restriction features only for individual browsers or users. The system allows you to extend these features to your entire network.

The system allows you to enforce:

- *Safe Search*—Supported in many major search engines, this service filters out explicit and adult-oriented content that business, government, and education environments classify as objectionable. The system does not restrict a user's ability to access the home pages for supported search engines.

You can use two methods to configure the system to enforce these features:

Method: Access Control Rules
: Content restriction features communicate the restricted status of a search or content query via an element in the request URI, an associated cookie, or a custom HTTP header element. You can configure access control rules to modify these elements as the system processes traffic.

Method: DNS Sinkhole
: For Google searches, you can configure the system to redirect traffic to the Google SafeSearch Virtual IP Address (VIP), which imposes filters for Safe Search.

The table below describes the differences between these enforcement methods.

Table 107: Comparison of Content Restriction Methods

Attribute	Method: Access Control Rules	Method: DNS Sinkhole
Supported devices	Any	Secure Firewall Threat Defense only
Search engines supported	Any tagged `safesearch supported` in the **Applications** tab of the rule editor	Google only
YouTube Restricted Mode supported	Yes	Yes
SSL policy required	Yes	No
Hosts must be using IPv4	No	Yes
Connection event logging	Yes	Yes

When determining which method to use, consider the following limitations:

- The access control rules method requires an SSL policy, which impacts performance.
- The Google SafeSearch VIP supports IPv4 traffic only. If you configure a DNS sinkhole to manage Google searches, any hosts on the affected network must be using IPv4.

The system logs different values for the **Reason** field in connection events, depending on the method:

- Access Control Rules—`Content Restriction`
- DNS Sinkhole—`DNS Block`

Requirements and Prerequisites for Content Restriction

Model Support

Any, or as indicated in the procedure.

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Guidelines and Limitations for Content Restriction

- Safe search is supported by Snort 2 only.
- YouTube and Google do not support the YouTubeEDU feature that was implemented in access control rules. Please remove any access control rules that configure YouTubeEDU as they are not truly functional. You can also remove associated decryption rules.

Using Access Control Rules to Enforce Content Restriction

The following procedure explains how to configure access control rules to restrict content."

Note When safe search is enabled in an access control rule, inline normalization is enabled automatically.

Procedure

Step 1 Create a decryption policy.

Step 2 Add rules for handling Safe Search traffic:

- Choose **Decrypt - Resign** as the **Action** for the rules.
- In **Applications**, add selections to the **Selected Applications and Filters** list:
  - Safe Search—Add the `Category: search engine` filter.

Step 3 Set rule positions for the rules you added. Click and drag, or use the right-click menu to cut and paste.

Step 4 Create or edit an access control policy, and associate the decryption policy with the access control policy.

For more information, see Associating Other Policies with Access Control, on page 1762.

Step 5 In the access control policy, add rules for handling Safe Search traffic:

- Choose **Allow** as the **Action** for the rules.
- In **Applications**, click the icon for **Safe search** ( ) and set related options.
  - Safe Search Options for Access Control Rules, on page 1976
- In **Applications**, refine application selections in the **Selected Applications and Filters** list.

  In most cases, enabling Safe Search populates the **Selected Applications and Filters** list with the appropriate values. The system does not automatically populate the list if a Safe Search application is already present in the list when you enable the feature. If applications do not populate as expected, manually add them as follows:
  - Safe Search—Add the `Category: search engine` filter.

For more information, see Configuring Application Conditions and Filters, on page 1789.

Step 6 Set rule positions for the access control rules you added. Click and drag, or use the right-click menu to cut and paste.

Step 7 Configure the HTTP response page that the system displays when it blocks restricted content; see Choosing HTTP Response Pages, on page 1875.

Step 8 Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Safe Search Options for Access Control Rules

The system supports Safe Search filtering for specific search engines only. For a list of supported search engines, see applications tagged `safesearch supported` in the **Applications** tab of the access control rule editor. For a list of unsupported search engines, see applications tagged `safesearch unsupported`.

When enabling Safe Search for an access control rule, set the following parameters:

Enable Safe Search
Enables Safe Search filtering for traffic that matches this rule.

Unsupported Search Traffic
Specifies the action you want the system to take when it processes traffic from unsupported search engines. If you choose **Block** or **Block with Reset**, you must also configure the HTTP response page that the system displays when it blocks restricted content; see Choosing HTTP Response Pages, on page 1875.

Using a DNS Sinkhole to Enforce Content Restriction

Typically, a DNS sinkhole directs traffic away from a particular target. This procedure describes how to configure a DNS sinkhole to redirect traffic to the Google SafeSearch Virtual IP Address (VIP), which imposes content filters on Google and YouTube search results.

Because Google SafeSearch uses a single IPv4 address for the VIP, hosts must use IPv4 addressing.

Caution If your network includes proxy servers, this content restriction method is not effective unless you position your Firewall Threat Defense devices between the proxy servers and the Internet.

This procedure enforces content restriction for Google searches only. It does not apply to any other search engine.

Before you begin

This procedure applies to Firewall Threat Defense only, and requires the IPS license.

Procedure

Step 1 Obtain a list of supported Google domains via the following URL: https://www.google.com/supported_domains.

Step 2 Create a custom DNS list on your local computer, and add the following entries:

- To enforce Google SafeSearch, add an entry for each supported Google domain.
- To enforce YouTube Restricted Mode, add a "youtube.com" entry.

The custom DNS list must be in text file (.txt) format. Each line of the text file must specify an individual domain name, stripped of any leading periods. For example, the supported domain ".google.com" must appear as "google.com".

Step 3 Upload the custom DNS list to the Firewall Management Center; see Uploading New Security Intelligence Lists to the Secure Firewall Management Center, on page 1439.

Step 4 Determine the IPv4 address for the Google SafeSearch VIP. For example, run `nslookup` on forcesafesearch.google.com.

Step 5 Create a sinkhole object for the SafeSearch VIP; see Creating Sinkhole Objects, on page 1440.

Use the following values for this object:

- IPv4 Address—Enter the SafeSearch VIP address.
- IPv6 Address—Enter the IPv6 loopback address (`::1`).
- Log Connections to Sinkhole—Click Log Connections.
- Type—Choose **None**.

Step 6 Create a basic DNS policy; see Creating Basic DNS Policies, on page 1900.

Step 7 Add a DNS rule for the sinkhole; see Creating and Editing DNS Rules, on page 1902.

For this rule:

- Check the **Enabled** check box.
- Choose `Sinkhole` from the **Action** drop-down list.
- Choose the sinkhole object you created from the **Sinkhole** drop-down list.
- Add the custom DNS list you created to the **Selected Items** list on **DNS**.
- (Optional) Choose a network in **Networks** to limit content restriction to specific users. For example, if you want to limit content restriction to student users, assign students to a different subnet than faculty, and specify that subnet in this rule.

Step 8 Associate the DNS policy with an access control policy; see Associating Other Policies with Access Control, on page 1762.

Step 9 Deploy configuration changes; see Deploy Configuration Changes, on page 259.

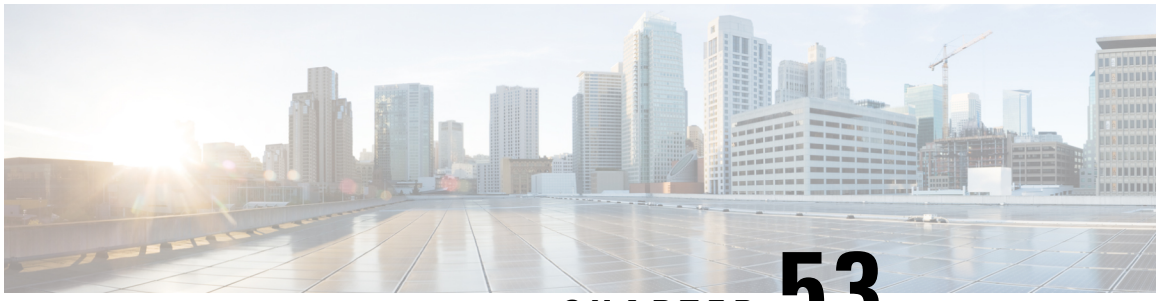

CHAPTER 53

Zero Trust Access

The following topics provide more information about Clientless Zero Trust Access, Universal Zero Trust Network Access; and how to configure and deploy them.

- About Clientless and Universal Zero Trust Network Access, on page 1979
- Clientless Zero Trust Access, on page 1979
- Overview of Universal Zero Trust Network Access, on page 1994
- History for Zero Trust Access, on page 2007

About Clientless and Universal Zero Trust Network Access

Zero Trust Access (ZTNA) is a general name for process that protect internal resources with identity-based access (including user trust and posture), and can provide intrusion protection, file, and malware protection.

Zero Trust Access is further divided into the following:

- *Clientless* ZTNA enables you to authenticate and authorize access to protected web based resources and applications from inside (on-premises) or outside (remote) the network using an external Security Assertion Markup Language (SAML) identity provider (IdP) policy. Use clientless ZTNA for remote users.
- *Universal* ZTNA provides identity-based network protection and can be used with both remote and on-premises users. On-premises users access a trusted network using a Secure Firewall Threat Defense while remote users install an application like Cisco Secure Client.

Clientless Zero Trust Access

Clientless Zero Trust AccessZero Trust Access is based on Zero Trust Network Access (ZTNA) principles. ZTNA is a zero trust security model that eliminates implicit trust. The model grants the least privilege access after verifying the user, the context of the request, and after analyzing the risk if access is granted.

Clientless Zero Trust AccessZero Trust Access enables you to authenticate and authorize access to protected web based resources and applications from inside (on-premises) or outside (remote) the network using an external Security Assertion Markup Language (SAML) identity provider (IdP) policy.

The features are:

- Supports multiple SAML-based identity providers such as Duo, Azure Active Directory (Azure AD), Okta, and other identity providers.

- Client applications such as Cisco Secure Client are not required on the endpoint (client devices) for secure access.
- Access and authentication is through the browser.
- Supports only TLS web applications.
- Client device posture is supported through agents such as Duo Health, using which the posture of the device can be evaluated against a policy in Duo, and access can be provided based on the same. The same functionality can be performed in conjunction with third-party identity providers that support posture evaluation with their agents such as Okta or PingID.
- Supports HTTP-Redirect SAML binding.
- Supports application groups that make it easier to enable zero trust accessclientless zero trust access protection on a set of applications.
- Leverages threat defense intrusion and malware protection on zero trust application traffic.

You can use the Secure Firewall Management Center web interface to create a Zero Trust Application Policy that allows you to define private applications and assign threat policies to them. The policy is application specific where the administrator decides the inspection levels based on the threat perception for that application.

How Threat Defense Works with Zero Trust Access

Figure 441: Threat Defense Deployment

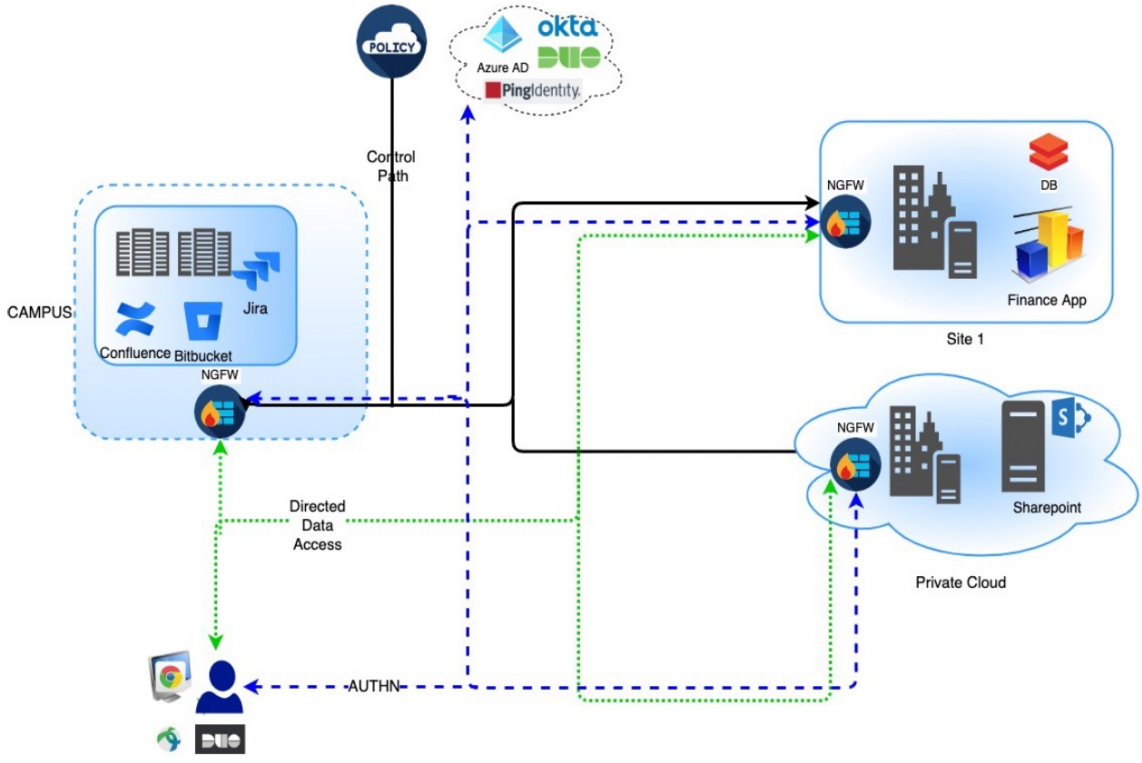

1. Using a browser, a remote or on-prem user sends a HTTPS request to connect to an application from an endpoint.
2. The HTTPS request is intercepted by the firewall that protects the application.
3. The firewall redirects the user to application's configured IdP for authentication.

Note In the figure, each firewall protects a set of web applications. The user can directly access the applications behind the firewall after authentication and authorization.

4. After the authentication and authorization process is complete, the firewall allows the user to access the application.

Why Use Zero Trust Access?

Zero Trust Access leverages the existing deployment of threat defense as an enforcement point to application access. It allows for segmented access to a private application with per application authorization and per application tunnel for remote and on-premises users.

The feature hides the network from users and allows users to only access applications they are authorized for. Authorization for one application in the network does not give an implicit authorization for other applications on the network, thereby reducing the attack surface significantly. In other words, every access to an application must be explicitly authorized.

Adding the zero trust access functionality to threat defense enables migration to a more secure access model without having to install or manage yet another device in the network.

The feature is easy to manage as it does not require a client and is per application access.

Components of a Zero Trust Access Configuration

A new configuration consists of a Zero Trust Application Policy, Application Group, and Applications.

- **Zero Trust Application Policy**— Consists of application groups, and grouped or ungrouped applications. Security Zones and Security Controls settings are associated at a global level for all the ungrouped applications and group of applications.

  A global port pool is assigned to the policy, by default. A unique port is automatically assigned from this pool to each private application that is configured.

  Zero Trust Application policy consists of application groups, and grouped or ungrouped applications.

- **Application Groups**—Consists of a logical group of applications that share SAML authentication settings and can optionally share Security Zones and Security Controls settings.

  Application Groups inherit the Security Zones and Security Controls settings from the global policy and can override the values.

  When an Application Group is created, the same SAML IdP configuration can be used for authenticating multiple applications. Applications that are part of an Application Group inherit the Application Group's SAML configuration. This eliminates the need to configure the SAML settings for each application. After the Application Group is created, new applications can be added to it without configuring the IdP for it.

When an end user tries to access an Application that is part of group, the user is authenticated to the Application Group for the first time. When the user tries to access other applications that are part of the same Application Group, the user is provided access without being redirected again to the IdP for authentication. This prevents overloading the IdP with requests for application access and optimizes the usage of the IdP if a limit is enabled.

- **Applications**—There are two types:
  - **Ungrouped Applications**— Are standalone applications. SAML settings must be configured for every application. The applications inherit the Security Zones and Security Controls settings from the global policy and can be overridden by the application.
  - **Grouped Applications**— Are multiple applications that are grouped under an Application Group. The SAML settings are inherited from the Application Group and cannot be overridden. However, the Security Zones and Security Controls settings can be overridden for each application.

The following certificates are required for the configuration:

- **Identity Certificate**—This certificate is used by Firewall Threat Defense to masquerade as the applications.Firewall Threat Defense behaves as a SAML Service Provider (SP). This certificate must be a wildcard or Subject Alternative Name (SAN) certificate that matches the FQDN of the private applications. It is a common certificate for all applications protected by Firewall Threat Defense.
- **IdP Certificate**—The IdP provides a certificate for each defined Application or Application Group. This certificate must be configured so that Firewall Threat Defense can verify the IDP's signature on incoming SAML assertions.

Note IdP certificates are commonly included within the metadata file; otherwise, users are required to have the IdP certificate readily available during the configuration of applications.

- **Application Certificate**—The encrypted traffic from user to the application is decrypted by Firewall Threat Defense using this certificate for the purpose of inspection.

Note This certificate is required to verify the cookies in the header to authorize connections, even if we are not conducting an IPS/Malware inspection.

Zero Trust Access Workflow

This figure depicts the Zero Trust Access workflow.

Figure 442: Zero Trust Access Workflow

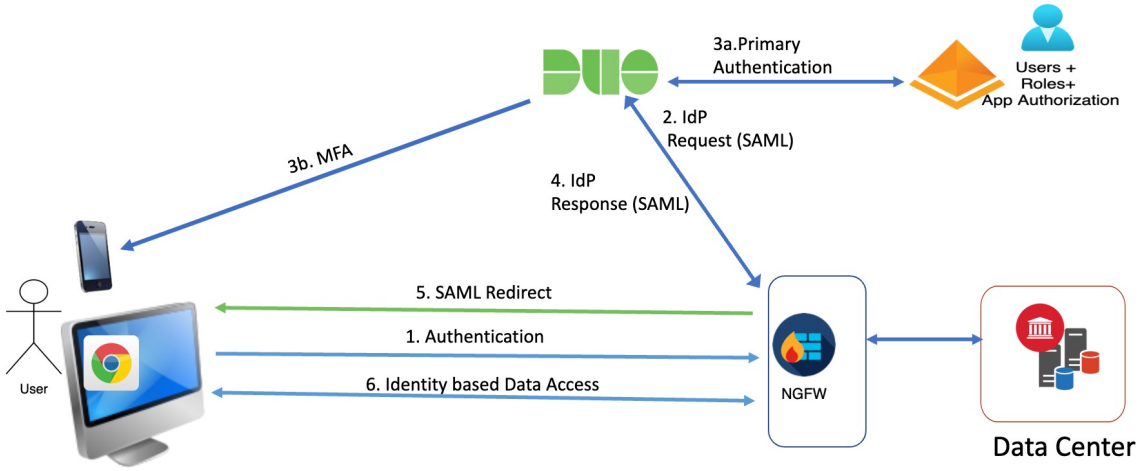

The workflow is as follows:

1. User types the application URL in the browser.
   - If the HTTPS request is valid, the user is redirected to the mapped port (Step 6).
   - If the HTTPS request is invalid, the user is sent for authentication per application (Step 2).
2. The user is redirected to the configured identity provider (IdP).
3. **a.** The user is redirected to the configured primary authentication source.
   b. The user is challenged with the configured secondary multi-factor authentication, if any.
4. The IdP sends a SAML response to threat defense. The user ID and other necessary parameters are retrieved from the SAML response through the browser.
5. The user is redirected to the application.
6. The user is allowed access to the application after validation is successful.

Limitations for Zero Trust Network Access

- Supports only TLS-enabled web applications. All traffic must be decrypted.
- Supports only SAML identity providers.
- Supports IPv6; however, you can use only homogeneous network address translations (NAT) scenarios such as NAT66 and NAT44.

  NAT64, and NAT46 scenarios are not supported.
- Zero trust network access is available on Firewall Threat Defense only if Snort 3 is enabled.
- By default, the Firewall Threat Defense device uses port 443 for secure communications. Because this is the same port used by TLS-enabled applications likely to be configured for zero trust network access, you must change the device's HTTP server port.

Go to **Devices** > **Platform Settings**, then edit the Threat Defense Settings policy for the device, click **HTTP Access**, select the **Enable HTTP Server** check box, and enter a port other than 443 in the field. When you're finished, click **Save** in the upper right corner of the page.

The following figure shows an example.

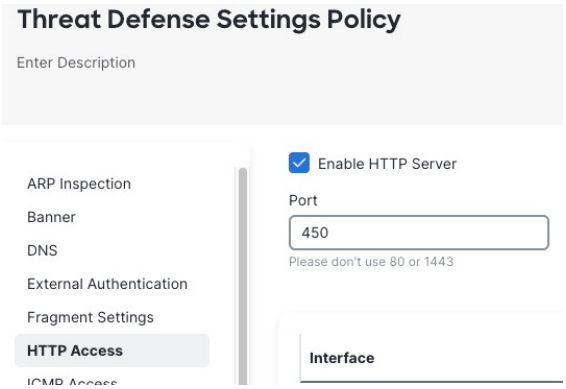

- All hyperlinks in protected web applications must have a relative path and are not supported on individual mode clusters.
- Protected web applications running on a virtual host or behind internal load balancers must use the same external and internal URL.
- Not supported on applications with strict HTTP Host Header validation enabled.
- If the application server hosts multiple applications and serves content based on the Server Name Indication (SNI) header in the TLS Client Hello, the external URL of the zero trust application configuration must match the SNI of that specific application.

Prerequisites for Zero Trust Application Policy

Prerequisite Type	Description
Licensing	• Smart license account with export-controlled features • (Optional) IPS and Threat licenses—It is required if security controls are used.
Configurations	Create a wildcard or Subject Alternative Name (SAN) certificate that matches the FQDN of private applications. For more information, see Adding Certificate Enrollment Objects, on page 1414.
	Create a security zone through which access to private applications are regulated. For more information, see Create Security Zone and Interface Group Objects, on page 777.

Manage Zero Trust Application Policies

You can create, edit, and delete zero trust application policies.

Procedure

Step 1 Choose **Policies** > **Zero Trust Application**

Step 2 Manage the zero trust access policies:

- Create—Click **New Policy**. See Create a Zero Trust Application Policy, on page 1985
- Edit—Click **Edit** (). See Edit a Zero Trust Application Policy, on page 1990
- Report—Click **Report** ().
- Delete—Click **Delete** ().

Step 3 Click **Save**.

What to do next

Ensure that there are no warnings before you deploy the configuration to threat defense. To deploy configuration changes, see Deploy Configuration Changes in the Cisco Secure Firewall Management Center Administration Guide.

Create a Zero Trust Application Policy

This tasks configures a Zero Trust Application Policy.

Before you begin

Ensure that you complete all the prerequisites listed in Prerequisites for Zero Trust Application Policy, on page 1984.

Procedure

Step 1 Choose **Policies** > **Zero Trust Application**.

Step 2 Click **Add Policy**.

Step 3 In the **General** section, enter the policy name in the **Name** field. The description field is optional.

Step 4 Enter a domain name in the **Domain Name** field.

Ensure that the domain name is added to the DNS. The domain name resolves to the Firewall Threat Defense gateway interface from where the application is accessed. The domain name is used to generate the ACS URL for all private applications in an Application Group.

Step 5 Choose an existing certificate from the **Identity Certificate** drop-down list.

Click the **Add** (+) icon to configure a certificate enrollment object. For more information, see Adding Certificate Enrollment Objects, on page 1414.

Step 6 Choose a security zone from the **Security Zones** drop-down list.

Click the **Add** (+) icon to add a new security zone.

To add security zones, see Create Security Zone and Interface Group Objects, on page 777.

Step 7 In the **Global Port Pool** section, a default port range is displayed. Modify, if required. Port values range from 1024 to 65535. A unique port from this pool is assigned to each private application.

Note
This port range should avoid any conflicts with the existing NAT range.

Step 8 (Optional) In the **Security Controls** section, you can add an Intrusion or Malware and File policy:

- **Intrusion Policy**—Choose a default policy from the drop-down list or click the **Add** (+) icon to create a new custom intrusion policy. For more information, see Creating a Custom Snort 3 Intrusion Policy topic in the latest version of the Cisco Secure Firewall Management Center Snort 3 Configuration Guide.
- **Variable Set**—Choose a default variable set from the drop-down list or click the **Add** (+) icon to create a new variable set. For more information, see Creating Variable Sets, on page 1457.

  Note
  To use variable sets, you must have the Secure Firewall Threat Defense IPS license for your managed devices.

- **Malware and File Policy**—Choose an existing policy from the drop-down list. Click the **Add** (+) icon to create a new malware and file policy. For more information, see Managing File Policies, on page 2232.

Step 9 Click **Save** to save the policy.

What to do next

1. Create an Application Group. See Create an Application Group, on page 1986.
2. Create an Application. See Create an Application, on page 1988.
3. Associate a Zero Trust Application Policy with a device. See Set Targeted Devices for Zero Trust Access Policy, on page 1989
4. Deploy configuration changes. See Deploy Configuration Changes in the Cisco Secure Firewall Management Center Administration Guide.

Create an Application Group

Before you begin

Create a Zero Trust Application Policy, on page 1985

Procedure

Step 1 Click **Add Application Group**.

Step 2 In the **Application Group** section, type the name in the **Name** field and click **Next**.

Step 3 In the **SAML Service Provider (SP) Metadata** section, the data is dynamically generated. Copy the values of the **Entity ID** and **Assertion Consumer Service (ACS) URL** fields or click **Download SP Metadata** to download this data in XML format for adding it to the IdP. Click **Next**.

Step 4 In the **SAML Identity Provider (IdP) Metadata** section, add the metadata using any one of the methods:

- **XML File Upload**—Choose a file or drag and drop the XML file.

  The details of the **Entity ID**, **Single Sign-On URL**, and **IdP Certificate** are displayed.

- **Manual Configuration**—Perform these steps:
  - **Entity ID**—Enter the URL that is defined in the SAML IdP to identify a service provider uniquely.
  - **Single Sign-On URL**—Enter the URL for signing into the SAML identity provider server.
  - **IdP Certificate**—Choose the certificate of the IdP enrolled in threat defense to verify the messages signed by the IdP.

    Click the **Add** (+) icon to configure a new certificate enrollment object. For more information, see Add Certificate Enrollment, on page 1416.

- **Configure Later**—In the event you do not have the IdP metadata, you can configure it later.

Click **Next**.

Step 5 In the **Re-authentication Interval** section, enter the value in the **Timeout Interval** field and click **Next**.

The re-authentication interval allows you to provide a value that determines when a user must authenticate again.

Step 6 In the **Security Zones and Security Controls** section, the security zones and threat settings are inherited from the parent policy. You can override these settings. Click **Next**.

Step 7 Review the configuration summary. Click **Edit** to modify the details in any of the sections. Click **Finish**.

Step 8 Click **Save**.

The Application Group is created and is displayed on the Zero Trust Application page.

What to do next

1. Create an Application, on page 1988.
2. Deploy configuration changes. See Deploy Configuration Changes in the Cisco Secure Firewall Management Center Administration Guide.

Create an Application

Use this task to create a Grouped or Ungrouped Application.

Before you begin

1. Create a Zero Trust Application Policy, on page 1985.
2. Create an Application Group, on page 1986 (required only for Grouped Applications).

Procedure

Step 1 Choose **Policies** > **Zero Trust Application**

Step 2 Choose the policy.

Step 3 Click **Add Application**.

Step 4 In the **Application Settings** section, complete the following fields.

- **Application Name**—Enter the application name.
- **External URL**—Enter the URL that is used by the user to access the application.
- **Application URL**—By default, the external URL is used as the Application URL. Uncheck the **Use External URL as Application URL** check box to specify a different URL.

  If Threat Defense uses an internal DNS, the Application URL must match an entry within that DNS to ensure it resolves to the application.
- **Application Certificate**—Choose the certificate for the private application. Click the **Add** (+) icon to configure an internal certificate object. For more information, see Adding Internal Certificate Objects, on page 1412.
- **IPv4 Source Translation**—Choose the source network for NAT from the list. Click **Add** (+) to create a network object. For more information, see Network, on page 1401.

  This Network Object or Object Group is used to translate a public network source IPv4 address of an incoming request to a routable IPv4 address inside the corporate network.

  Note
  Only object or object groups of type Host or Range are supported.
- **Application Group**—Choose the Application Group from the drop-down list. See Create an Application Group, on page 1986.

  Note
  This field is not applicable for an ungrouped application.

Step 5 Click **Next**.

Step 6 Depending on the type of application:

- For a Grouped Application, the **SAML Service Provider (SP) Metadata**, **SAML Identity Provider (IdP) Metadata** , and **Re-authentication Interval** are inherited from the Application Group and do not need to be configured by the user.

- For an Ungrouped Application, perform these steps:

  a. In the **SAML Service Provider (SP) Metadata** section, the data is dynamically generated. Copy the **Entity ID**, **Assertion Consumer Service (ACS) URL** of the IdP, or click **Download SP Metadata** to download this data in XML format for adding it to the IdP. Click **Next**.

  b. In the **SAML Identity Provider (IdP) Metadata** section, add the metadata using any one of the methods:

  - **XML File Upload**—Choose a file or drag and drop the XML file.

    The details of the **Entity ID**, **Single Sign-On URL**, and **IdP Certificate** are displayed.

  - **Manual Configuration**—Perform these steps:

    - **Entity ID**—Enter the URL that is defined in the SAML IdP to identify a service provider uniquely.
    - **Single Sign-On URL**—Enter the URL for signing into the SAML identity provider server.
    - **IdP Certificate**—Choose the certificate of the IdP enrolled in threat defense to verify the messages signed by the IdP.

      Click the **Add** (+) icon to configure a new certificate enrollment object. For more information, see Add Certificate Enrollment, on page 1416.

  - **Configure Later**—If you do not have the IdP metadata, you can configure it later.

  Click **Next**.

  c. In the **Re-authentication Interval** section, enter the value in the **Timeout Interval** field and click **Next**. The reauthentication interval allows you to provide a value that determines when a user must authenticate again.

Step 7 In the **Security Zones and Security Controls** section, the security zones and threat settings are inherited from the parent policy or application group. You can override these settings. Click **Next**.

Step 8 Review the configuration summary. Click **Edit** to modify the details in any of the sections. Click **Finish**.

Step 9 Click **Save**.

The application is listed on the Zero Trust Application page and enabled by default.

What to do next

1. Set Targeted Devices for Zero Trust Access Policy, on page 1989.
2. Deploy configuration changes. See Deploy Configuration Changes in the Cisco Secure Firewall Management Center Administration Guide.

Set Targeted Devices for Zero Trust Access Policy

Each Zero Trust Application policy can target multiple devices; each device can have one deployed policy at a time.

Before you begin

1. Create a Zero Trust Application Policy, on page 1985.
2. Create an Application Group, on page 1986.
3. Create an Application, on page 1988.

Procedure

Step 1 Choose **Policies** > **Zero Trust Application**

Step 2 Choose the policy.

Step 3 Click **Targeted Devices**.

Step 4 Choose the devices where you want to deploy the policy using any one of the methods:

- Choose a device in the **Available Devices** list and click >> or the **Add** ( ) icon.
- To remove a device from the **Selected Devices** list, choose a device and click << or the **Delete** ( ) icon.

Step 5 Click **Apply** to save policy assignments.

Step 6 Click **Save** to save the policy.

What to do next

Deploy configuration changes. See Deploy Configuration Changes in the Cisco Secure Firewall Management Center Administration Guide.

Edit a Zero Trust Application Policy

You can edit the settings of a Zero Trust Application Policy, Application Group, or Application.

Procedure

Step 1 Choose **Policies** > **Zero Trust Application**

Step 2 Click **Edit** (✎) next to the Zero Trust Application Policy you want to edit.

Step 3 Edit your Zero Trust Application Policy.

You can change the following settings or perform these actions:

- Name and Description—Click **Edit** (✎) next to the policy name, make your changes, and click **Apply**.
- To modify the policy settings:
  - Click **Settings**

- Modify the settings as required.

  Important
  Editing the domain name for the SAML ACS URL interrupts application access.

- Click **Save**.

- To modify the Application Group settings:
  - Click **Applications**.
  - Click **Edit** (✎) next to the Application Group you want to edit.
  - In each section, click **Edit** to modify the settings, as required

    Important
    Editing the Application Group name interrupts application access.

  - Click **Apply** after you modify the settings in a section.
  - Click **Finish**.
  - Click **Save**.

- To modify the Application settings:
  - Click **Applications**.
  - Click **Edit** (✎) next to the Application you want to edit.
  - In each section, click **Edit** to modify the settings, as required.

    Important
    Editing the Application name interrupts application access.

  - Click **Apply** after you modify the settings in a section.
  - Click **Finish**.
  - Click **Save**.

- To enable, disable, or delete multiple Applications, choose the Applications, click the required bulk action, and click **Save**.

  Note
  These actions are also available in the right-click menu.

  - To enable all Applications, click **Bulk Actions** > **Enable**.
  - To disable all Applications, click **Bulk Actions** > **Disable**.
  - To delete all Applications, click **Bulk Actions** > **Delete**.

- Click **Return to Zero Trust Application** to return to the policy page.

What to do next

Deploy configuration changes. See Deploy Configuration Changes in the Cisco Secure Firewall Management Center Administration Guide.

Monitor Zero Trust Sessions

Connection Events

After a Zero Trust Application Policy is deployed, new fields are available. To add the fields to the table view:

1. Choose **Analysis** > **Connections**.
2. Click the **Table View of Connection Events** tab.
3. In the table view of events, multiple fields are hidden by default. To change the fields that are displayed, click the **x** icon in any column name to display a field selector.
4. Select the following fields:
   - **Authentication Source**
   - **Zero Trust Application**
   - **Zero Trust Application Group**
   - **Zero Trust Application Host**
   - **Zero Trust Application Policy**
   - **Zero Trust Origin User**
   - **Zero Trust Proxy**
   - **Zero Trust Rule**
   - **Zero Trust Status**
   - **Zero Trust Tunnel ID**
5. Click **Apply**.

See Connection and Security-Related Connection Events in the Secure Firewall Management Center Administration Guide for more information on the connection events.

Zero Trust Dashboard

The Zero Trust dashboard allows you to monitor real-time data from active zero trust sessions on the devices.

The Zero Trust dashboard provides a summary of the top zero trust applications and zero trust users that are managed by the management center. Choose **Overview** > **Dashboards** > **Zero Trust** to access the dashboard.

The dashboard has the following widgets:

- Top Zero Trust Applications
- Top Zero Trust Users

CLI Commands

Log in to the device CLI and use the following commands:

CLI Command	Description
show running-config zero-trust	To view the running configuration for a zero trust configuration.
show running-config zero-trust-hybrid	To view the running configuration for a universal zero trust configuration.
show zero-trust	To display the run-time zero trust statistics and session information.
show cluster zero-trust	To display the summary of zero trust statistics across nodes in a cluster.
clear zero-trust	To clear zero trust sessions and statistics.
show counters protocol zero_trust	To view the counters that are hit for zero trust flow.

Diagnostics Tool

The diagnostics tool facilitates the troubleshooting process by detecting possible issues with zero trust configurations. The diagnostics can be classified into two types:

- **Application-specific diagnostics** are used to detect issues such as:
  - DNS-related issues
  - Misconfigurations such as socket not open, and issues with classification and NAT rules.
  - Issues with deployment of zero trust policy or SSL rules
  - Issues with source NAT issues and exhaustion of PAT pool
- **General diagnostics** are used to detect issues such as:
  - Strong cipher license not enabled
  - Invalid application certificate
  - SAML-related issues
  - Home agent and cluster bulk sync issues

To run the diagnostic tool:

1. Click Diagnostics ( ) next to the zero trust application that you want to troubleshoot. The **Diagnostics** dialog box appears.
2. Choose the device from the **Select Device** drop-down list and click **Run**. A report is generated in the **Reports** tab after the diagnostic process is complete.
3. To view, copy, or download the logs, click the **Logs** tab.

Overview of Universal Zero Trust Network Access

Universal Zero Trust Network Access (universal ZTNA) enables administrators to specifically allow access to internal network resources according to user identity including user trust and posture, without granting access to the entire network as with Remote Access VPN. Universal ZTNA is a client-based ZTNA solution that enables users to securely access internal resources and applications regardless of their location, whether remote or on-premises.

Because universal ZTNA does not assume that access granted to one application implicitly authorizes access to other applications, the network attack surface is reduced.

Universal ZTNA ensures least-privileged, per-application, per-user access with strong authentication, posture validation, and comprehensive traffic inspection. It secures applications effectively across hybrid environments.

Components of Universal ZTNA

A new configuration of universal ZTNA consists of Security Cloud Control Firewall Management (formerly called Cisco Defense Orchestrator), and Secure Access, both provisioned on the Security Cloud Control platform. Security Cloud Control Firewall Management manages the Firewall Threat Defense devices through the Secure Firewall Management Center.

Figure 443: Components of Universal ZTNA

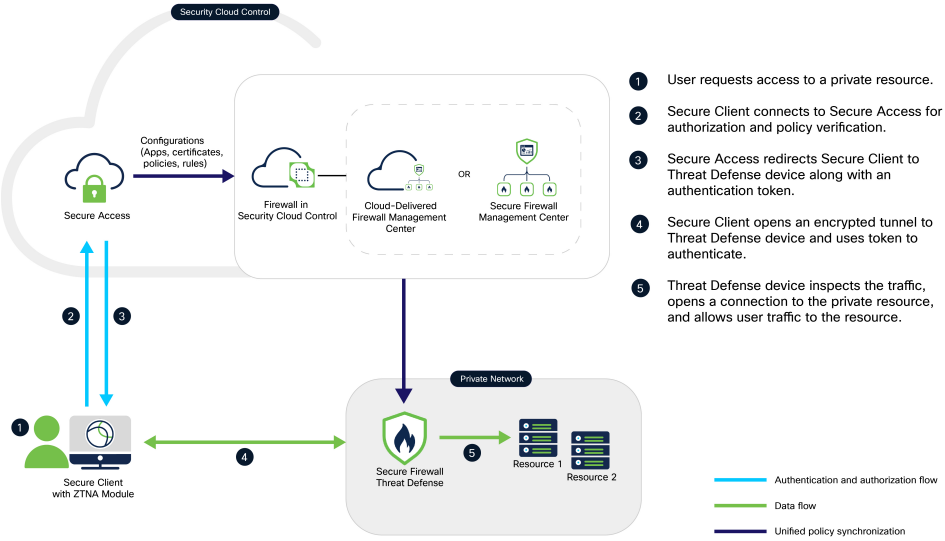

- **Security Cloud Control Firewall Management**: Manages the configuration and deployment of universal ZTNA policies to the Firewall Threat Defense devices. The Threat Defense devices protect on-premises resources by enforcing universal ZTNA policies. Threat Defense inspects traffic and enforces intrusion prevention system (IPS), file, and malware policies on the traffic.
- **Secure Access**: Secure Access defines the access policies, posture, and security profiles for the user. It enforces the policies for user traffic through the cloud.
- **Security Cloud Control platform**: Security Cloud Control provides a unified secure management plane for both Secure Access and Firewall, simplifying the administration of universal ZTNA policies across them.

- **Secure Client**: The Secure Client is installed on the end user's device. It acts as the enforcement point that intercepts connection requests to protected internal resources, enabling secure, identity-based access.

The following sections describe how to enable universal ZTNA on a Threat Defense device. For the complete configuration of universal ZTNA, refer to the *Universal Zero Trust Network Access Configuration Guide*.

How Threat Defense Works with Universal ZTNA

A universal ZTNA-enabled Firewall Threat Defense device is a critical enforcement point that protects private resources by integrating zero trust principles with firewall capabilities. Here's how it operates to secure private resources:

- **Enforces access and Validates Policy**: A Threat Defense device intercepts all access requests to private resources. It enforces granular, context-aware access policies that verify user identity, device posture, and contextual factors before granting access. Only requests that meet the least-privilege criteria are allowed, blocking the unauthorized attempts to access resources.
- **Establishes a Secure Tunnel to the Resource**: After validating the access policy, the Threat Defense device establishes a secure, encrypted tunnel between the user's device and the private resource. This tunnel ensures that private resources are isolated from direct network exposure.
- **Inspects Traffic**: Threat Defense device continuously inspects the traffic and blocks threats before they reach the resources. The device applies the intrusion prevention system, file, and malware detection capabilities to detect and block the threats.
- **Microsegments and Prevents Lateral Movement**: The device enforces microsegmentation by restricting traffic flows to only the authorized resources. This containment prevents lateral movement within the network, limiting the impact of any potential breach.

Prerequisites for Universal Zero Trust Network Access

This topic discusses requirements and guidelines for Universal Zero Trust Network Access (universal ZTNA).

Licensing Requirements

- Secure Firewall Management Center requires a smart license account with export-controlled features. It does not function in evaluation mode for universal ZTNA.

  Secure Firewall Threat Defense devices require an IPS license if Intrusion policy is configured. If Malware policies are configured, the devices require IPS and Malware Defense licenses. For more information, refer to the "Licenses" section in the *Cisco Secure Firewall Management Center Administration Guide*.

- Secure Access requires a subscription of Cisco Secure Private Access Essentials or Advantage.

Device Requirements

- All Secure Firewall Management Center and Secure Firewall Threat Defense devices must be running Version 7.7.10 or later.
- All Secure Firewall Threat Defense devices must be configured for routed mode; transparent mode is not supported.

- In Security Cloud Control, when you are configuring universal zero trust access for a device, ensure that the Enrollment Type for the device identity certificate is an object that is created using the PKCS12 file format. No other certificate type is supported. If necessary, you can also create a new certificate object from Security Cloud Control, which supports the PKCS12 format. See Configure Security Devices.
- Configure the Domain Name System (DNS) to resolve Fully Qualified Domain Name (FQDN) of private resources. Use the Platform Settings menu on the Secure Firewall to configure the DNS. See Interface and Device Settings.
- High Availability (HA) devices *are* supported; they are displayed as one entity.
- Secure Client (with ZTNA module enabled) Version 5.1.10 and later is supported.

  The client must be running in a platform that supports Trusted Platform Module (TPM), such as Windows 11.

Guidelines on Certificate Types

- **User Device Identity Certificate**: Secure Client, which is zero trust access enabled, presents the user identity certificate during the Mutual Transport Layer Security (mTLS) session with Secure Access and Firewall Threat Defense to request access to private resources.
- **Firewall Threat Defense Device Certificate**: Threat Defense devices that are universal ZTNA-enabled use device certificates to establish secure mTLS connections with the Secure Client and Secure Access. Ensure that the device identity certificate is of type PKCS12.

  If you have already enrolled a manual certificate for the device, first export it to the PKCS12 format using the **Devices** > **Certificates** > **Export Certificate** menu on Firewall Management Center. Use the exported PKCS12 file to create a new PKCS12 certificate enrollment object.
- **Decryption Certificate**: (Optional) To decrypt the traffic that is sent to private resources, enable **Decryption** for the resources in Secure Access and provide the server certificate and key. We recommend that you use a certificate that is signed by a publicly recognized certificate authority (CA).

Supported Devices

Both on-premises Firewall Management Center and cloud-delivered Firewall Management Center can be configured to manage the devices.

Only devices that have 16 cores or more are supported. Such models of Secure Firewall Threat Defense are:

- 1150
- 3105, 3110, 3120, 3130, 3140
- 4115, 4125, 4145, 4112
- 4215, 4225, 4245
- FTDv

Limitations of Universal Zero Trust Network Access

- Universal ZTNA does not support IPv6.

- Universal ZTNA-enabled devices do not enforce policies for traffic over a site-to-site tunnel.
- Universal ZTNA does not support clustered devices.
- Universal ZTNA sessions do not support jumbo frames.
- Currently, universal ZTNA supports only the United States and Europe regions.
- Universal ZTNA supports only global VRF.
- Universal ZTNA does not support protocols such as FTP or TFP, where the data or secondary connection originates from a server.

  For example, an active FTP connection uses a persistent control connection for commands and creates temporary data connections for file transfers. Universal ZTNA does not support such data connections that originate from the server.

Integrate Firewall Management Center with Security Cloud Control

Note This task is applicable only to on-premises Firewall Management Center.

Integrating the on-premises Secure Firewall Management Center with Cisco Security Cloud Control enables you to configure your Secure Firewall Management Center and its associated Secure Firewall Threat Defense devices. These devices can then use the networks, private resources, and policies necessary to configure and manage universal ZTNA.

Note Universal ZTNA uses *only* the access policies that are defined by Secure Access. Any other access control policies and rules deployed to the Threat Defense devices from the Secure Firewall Management Center are ignored for universal ZTNA.

Before you begin

Your Cisco contact must onboard your Cisco Security Cloud Control and Secure Access systems, and create users and tenants.

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Policies** > **Zero Trust Application**.

Step 3 Click the **Universal** tab.

The Zero Trust Access page appears.

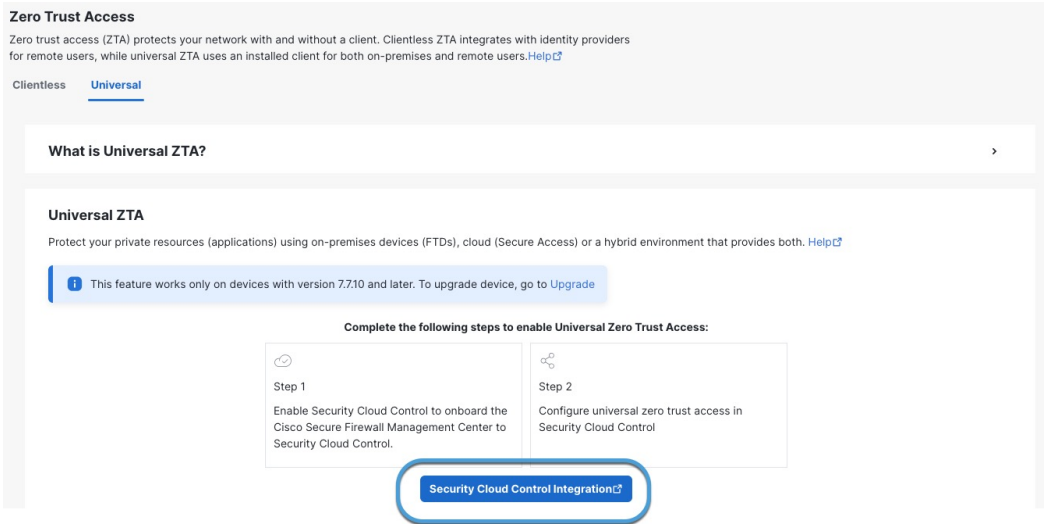

Step 4 Click **Security Cloud Control Integration**.

Step 5 From the **Current Cloud Region** list, click the name of your Cisco Security Cloud Control region.

Step 6 Click **Enable Cisco Security Cloud**.

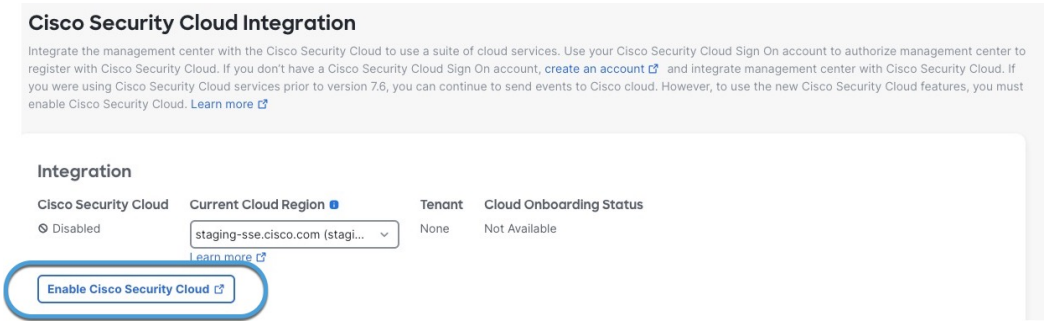

Step 7 When prompted, click **Continue to Cisco SSO**.

Step 8 Log in to Cisco Security Cloud Control.

Step 9 From the **Select Tenant** list, click the name of your tenant.

Step 10 At the following page, click **Authorize FMC**.

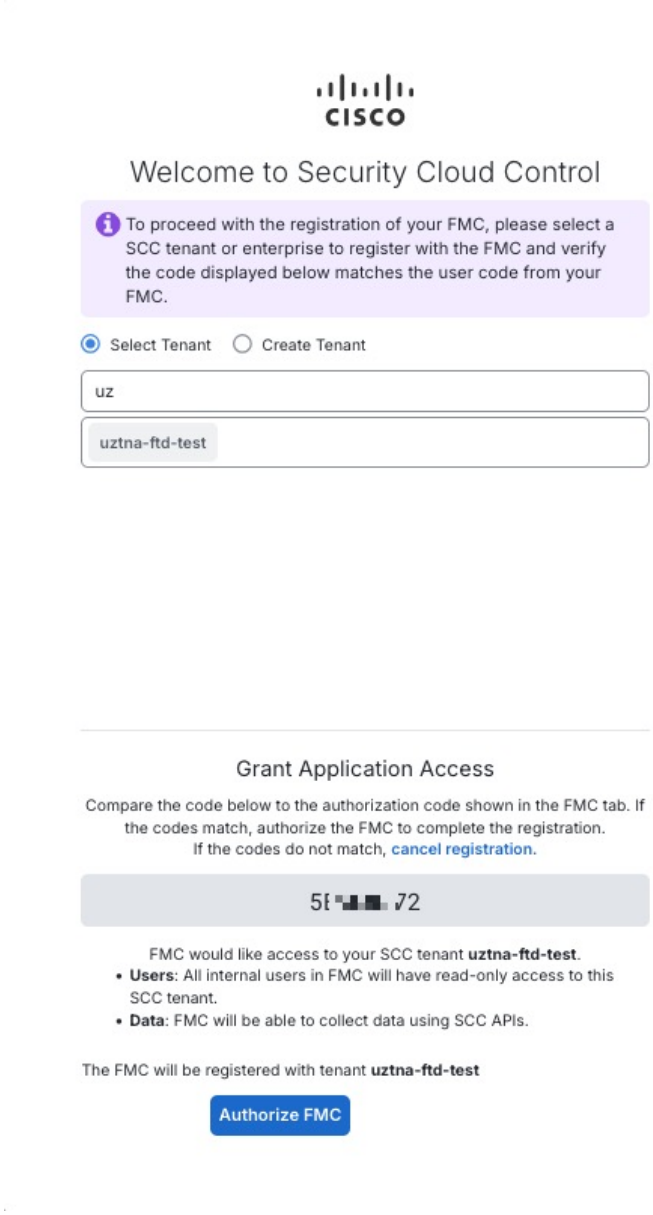

Step 11 When prompted, close the tab page.

Step 12 A confirmation message appears to indicate that the onboarding was successful.

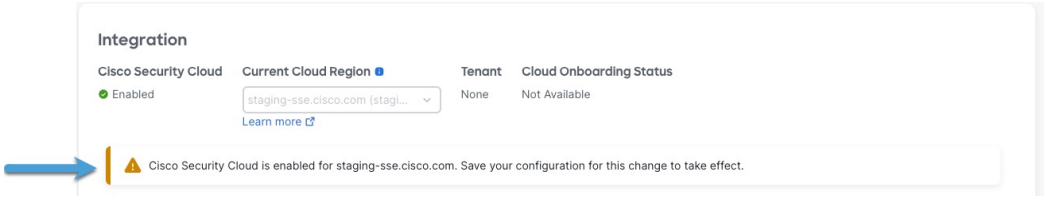

Step 13 Click **Save** at the bottom of the page.

It can take several minutes to save the configuration. After the configuration is saved, the page displays the onboarding status and the tenant name.

Step 14 For more information about other options on this page, see Security Cloud Control Settings, on page 2000.

Security Cloud Control Settings

The following topics discuss settings on the Cisco Security Cloud Integration page, which can be reached by choosing **Integration** > **Cisco Security Cloud** on Secure Firewall Management Center.

Event Configuration

Monitor the selected events in Cisco Security Cloud Control:

- **Send events to the cloud**: Select this check box to monitor events in Cisco Security Cloud Control; clear the check box to not monitor any events.

  You can view events in Cisco Security Cloud Control at **Firewall** > **Events & Logs**.

- **Intrusion events**: If you are using IPS policies, select this check box to monitor those policies.
- **File and malware events**: If you are using file or malware policies, select this check box to monitor those policies.
- **Connection events**: Select the check box next to the events to monitor, either **Security** (for file and IPS events) or **All**.

For more information about events, see Event Types in Security Cloud Control.

Cisco Security Cloud Support (Optional)

Optionally, select the check box to enable the following:

- **Enable Cisco Success Network**: Select the check box to enable collection of statistics discussed in *Cisco Success Network Telemetry Data Collected from Cisco Secure Firewall Management Center, Version 7.7*.
- **Enable Cisco Support Diagnostics**: Select the check box to enable collection of statistics discussed in the Cisco Secure Firewall Management Center Administration Guide.

Cisco AI Assistant for Security

Select the check box to enable the AI assistant as discussed in Use Cisco AI Assistant for Security to Manage Your Threat Defense Devices Effectively.

Cisco XDR Automation

Select the check box to enable XDR workflow automation as discussed in the Cisco XDR Help Center.

Policy Analyzer and Optimizer

Select the check box to enable the policy analyzer; click **Learn more** for details.

Zero-Touch Provisioning (ZTP)

Select the check box to enable zero-touch provisioning as discussed in the Cisco Secure Firewall Management Center Administration Guide.

Configure Security Devices

All Firewall Threat Defense devices associated with the Secure Firewall Management Center that you onboarded to Cisco Security Cloud Control are *security devices* to which you can:

- Associate private resources, which are internal applications you want to protect with identity-based access control, IPS, malware, and other protections.
- Deploy Secure Access access rules. Security devices are responsible for enforcing access rules for on-premises users, remote users, or both.

Perform these steps to enable **universal zero trust network access settings** on the Threat Defense devices. These steps include configuring the device FQDN, inside interface, outside interface, and PKCS12 certificate to enable universal ZTNA on the devices.

Before you begin

Complete all tasks discussed in Integrate Firewall Management Center with Security Cloud Control, on page 1997.

You must know the name of each device's internal and external network interfaces:

- The internal interface (also referred to as the *DMZ* interface) is used to apply access rules to on-premises users.
- The external interface is used to apply access rules to remote users.

You can choose internal, external, or both types of interfaces for each security device.

Procedure

Step 1 In the Secure Firewall Management Center, click **Policies** > **Zero Trust Application**.

Step 2 Click **Configure Universal ZTA in Security Cloud Control**.

This figure shows an example.

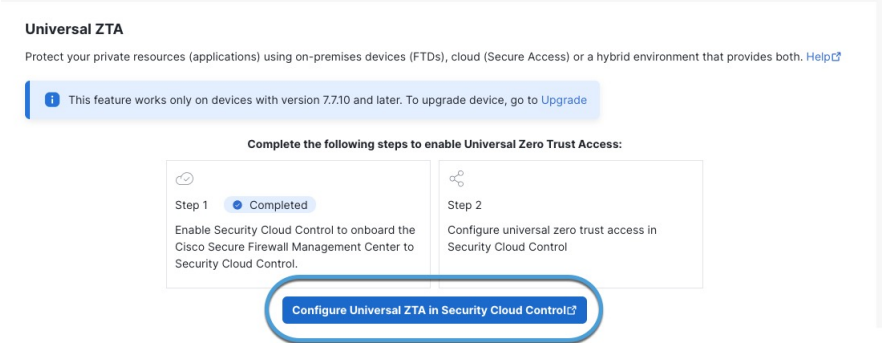

Step 3 When prompted, log in to Cisco Security Cloud Control.

Step 4 When prompted, select your organization from the drop-down list and click **Continue**.

Select an organization that has both Secure Access and Secure Firewall micro applications configured.

This figure shows an example.

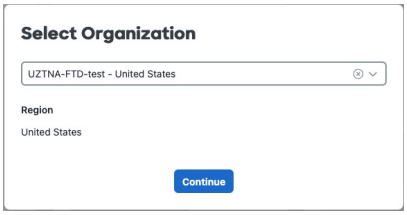

Step 5 In Cisco Security Cloud Control, in the Products section, click **Firewall**.

This figure shows an example.

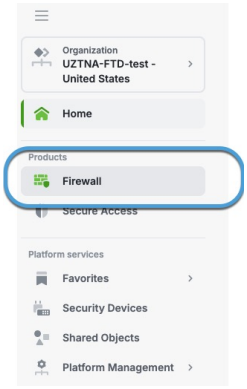

Step 6 In the Manage section, click **Security Devices**.

The **Security Devices** page displays the available security devices.

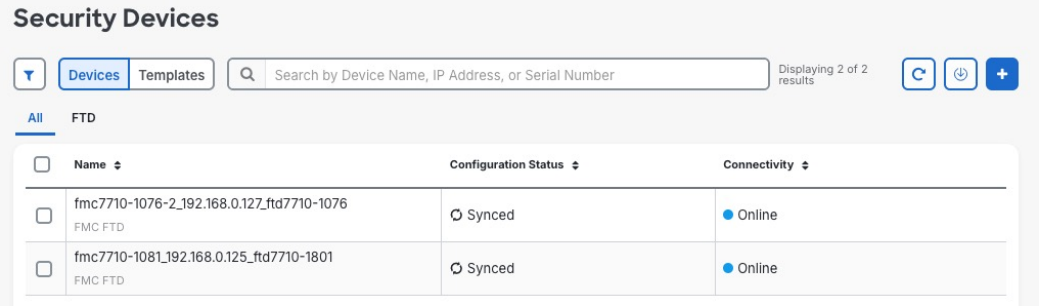

Step 7 Select the check box next to a device to add to the universal zero trust network access configuration.

Step 8 In the right pane, click **Device Management** > **Universal zero trust access settings**.

This figure shows an example.

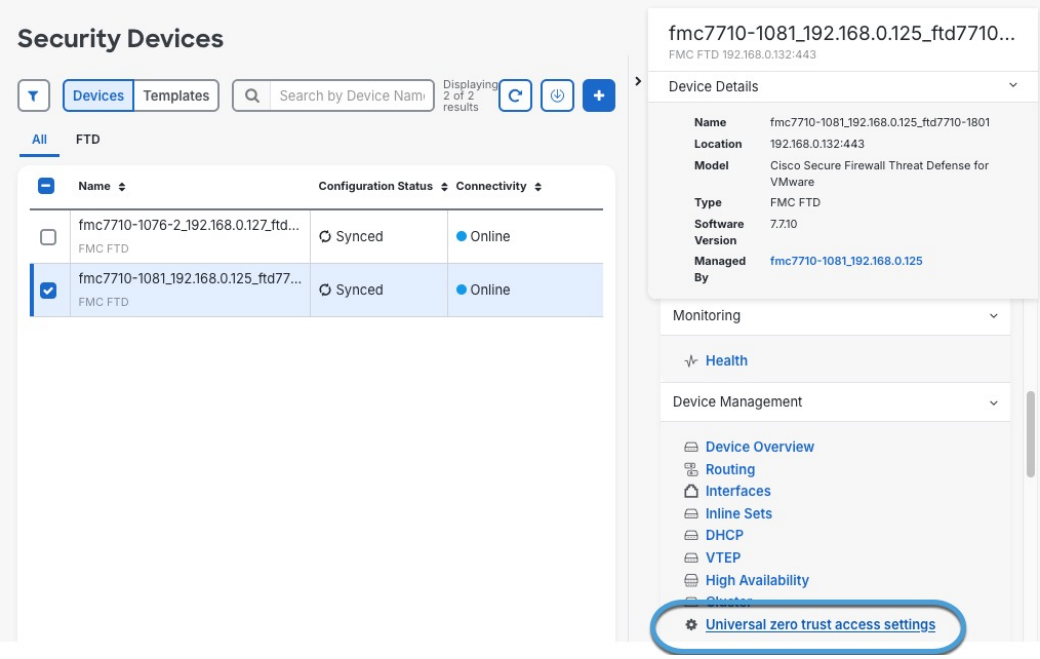

Step 9 Enter or edit the following information on the **Configure device for Universal Zero Trust Access** page.

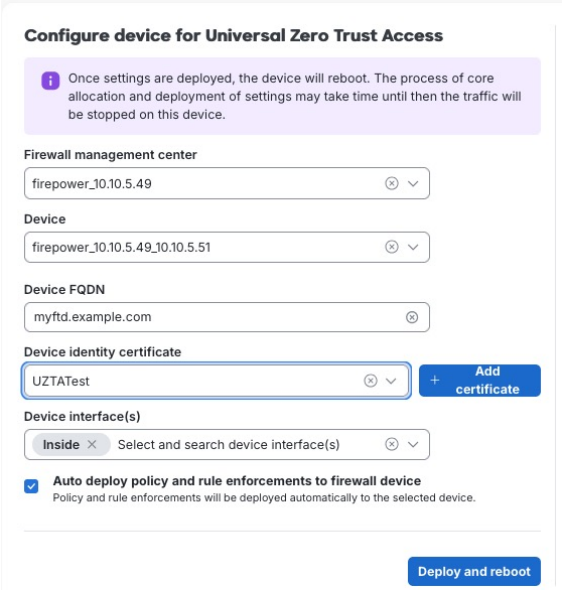

This table describes the configurations to enable universal ZTNA on the device.

Item	Description
Firewall management center	From the drop-down list, click the name of a Secure Firewall Management Center to use for policy deployment, monitoring, and other tasks.
Device	From the drop-down list, click the name of a device to use for rule deployment and enforcement.
Device FQDN	Enter the security device's fully qualified domain name (FQDN). The FQDN is also referred to as the TLS/SSL certificate's Common Name. The **Device identity certificate** must have a Common Name that either: • *Exactly matches* the value you enter in this field. • Matches a Subject Alternative Name (SAN) in the certificate. For more information, consult a resource such as What is the Common Name? on ssl.com.

Item	Description
Device identity certificate	From the drop-down list, click the name of an existing identity certificate from the list. Click **Add certificate** and add an identity certificate in **.p12** format (also referred to as PKCS#12; see this article on ssl.com). **Note** Universal ZTNA supports only the PKCS#12 format of certificate enrollment. In the provided fields, enter a **Name** to identify the certificate. Then copy/paste, drag/drop, or upload the certificate and private key. If the certificate is encrypted, enter its password in the provided field. You can optionally use a wildcard certificate as discussed in What is a Wildcard Certificate? on ssl.com.
Device Interface(s)	From the drop-down list, select the check box next to any of the following types of interfaces. • Internal network interface (or DMZ): deploys access rules for on-premises users only. • External network interface: deploys access rules for remote users only. • Both types of interfaces: deploys access rules for either on-premises or remote users.
Auto deploy policy and rule enforcements to firewall device	Select the check box to automatically deploy access rules to the device after they are updated on Secure Access. On the device, the Auto deploy feature selectively deploys only the Universal ZTNA access policy. It does not impact other changes or configurations on the Firewall Management Center. **Note** If there are other interdependent policies on the device (which are interlinked with the Universal ZTNA access policy), the Firewall configuration status displays an error message. The deployment then stops. In such cases, you should manually deploy the Universal ZTNA access policy from the Firewall Management Center.

Step 10 Click **Deploy and Reboot**.

The device reboots to reallocate the system resources for universal ZTNA components.

Note
The device takes several minutes to reboot, during which time all traffic handled by the device is disrupted.

If you deploy a High Availability (HA) pair of devices, both devices reboot simultaneously.

Step 11 On the **Security Devices** page, select the check box next to the device to which you just deployed the Universal ZTNA configuration.

The right pane displays the deployment status, as shown in the figure.

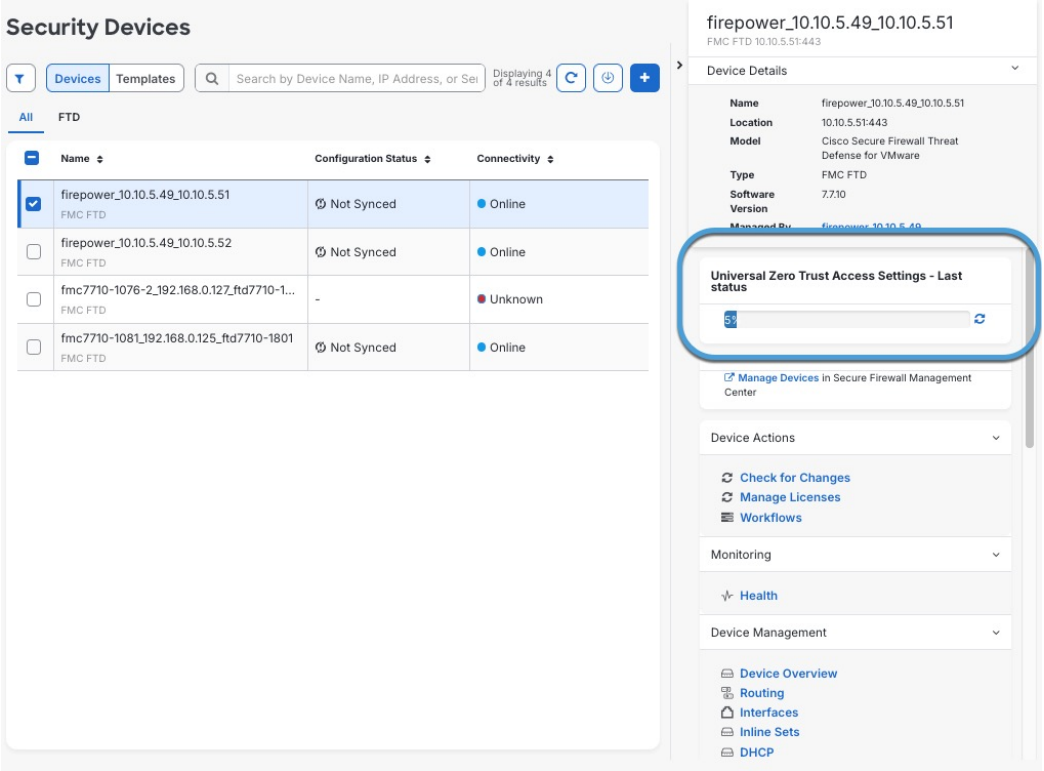

For additional information, click **Device Actions** > **Workflows** in the right pane.

After the deployment completes, you can view the completion status in the **Universal Zero trust Access Settings - Last status** tab for the device.

Universal ZTNA-enabled Firewall Threat Defense device is connected to Secure Access.

What to do next

Check the availability of the Threat Defense device under Secure Access by clicking **Security Cloud Control** > **Secure Access** > **Connect** > **Network Connections** > **FTD**.

History for Zero Trust Access

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Universal Zero Trust Network Access (universal ZTNA).	7.7.10	7.7.10	Universal Zero Trust Network Access (universal ZTNA) is a comprehensive solution that provides secure access to internal network resources based on user identity, trust, and posture. It ensures that access to one application does not implicitly grant access to the entire network, as with remote access VPN. New/modified screens: **Policies** > **Zero Trust Application** Requires Cisco Secure Access and Security Cloud Control. Deployment restrictions: Not supported with clustered devices, container instances, or transparent mode. Supported platforms: Secure Firewall 1150, 3100, 4100, 4200, and Firewall Threat Defense Virtual.
Zero trust access enhancements.	7.4.1	7.4.1 with Snort 3	Management center now includes the following zero trust access enhancements: • You can configure source NAT for an application. The configured network object or object group translates the incoming request's public network source IP address to a routable IP address inside the application network. • You can troubleshoot the zero trust configuration issues using the diagnostics tool. New/modified screens: **Policies** > **Access Control** > **Zero Trust Application** New/modified CLI commands: **show running-config zero-trust**, **show zero-trust statistics**

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Clientless zero-trust access.	7.4.0	7.4.0 with Snort 3	Zero Trust Access allows you to authenticate and authorize access to protected web based resources, applications, or data from inside (on-premises) or outside (remote) the network using an external SAML Identity Provider (IdP) policy. The configuration consists of a Zero Trust Application Policy (ZTAP), Application Group, and Applications. New/modified screens: • **Policies** > **Zero Trust Application** • **Analysis** > **Connections** > **Events** • **Overview** > **Dashboard** > **Zero Trust** New/modified CLI commands: • **show running-config zero-trust application** • **show running-config zero-trust application-group** • **show zero-trust sessions** • **show zero-trust statistics** • **show cluster zero-trust statistics** • **clear zero-trust sessions application** • **clear zero-trust sessions user** • **clear zero-trust statistics**

PART VIII

SD-WAN

- SD-WAN Capabilities, on page 2011

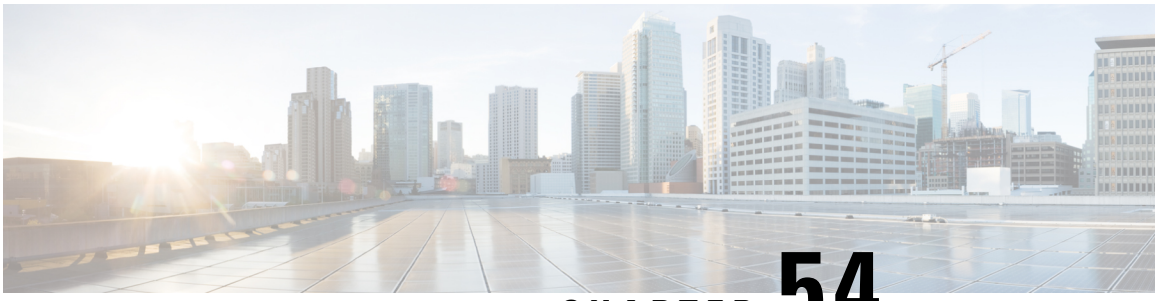

CHAPTER 54

SD-WAN Capabilities

This chapter describes the SD-WAN capabilities supported in the management center.

- Overview of SD-WAN Capabilities, on page 2011
- Features, on page 2012
- Using SD-WAN Wizard for Secure Branch Network Deployment, on page 2013
- Use Cases for SD-WAN Capabilities, on page 2028

Overview of SD-WAN Capabilities

Software-Defined WAN (SD-WAN) solutions replace traditional WAN routers and are agnostic to WAN transport technologies. SD-WAN provides dynamic, policy-based, application path selection across multiple WAN connections and supports service chaining for additional services such as WAN optimization and firewalls.

As organizations expand their operations across multiple branch locations, ensuring secure and streamlined connectivity becomes paramount. Deploying a secure branch network infrastructure involves complex configurations, which can be time-consuming and prone to configuration errors if not handled properly. However, organizations can overcome these challenges by leveraging the Cisco Secure Firewall Management Center (management center) and the Cisco Secure Firewall Threat Defense (threat defense) devices for a simplified and secure branch deployment.

In this guide, we explore the concept of simplifying secure branch deployment using a robust firewall solution. By integrating a secure firewall as a foundational component of the branch network architecture, organizations can establish a strong security baseline while simplifying the deployment process. This approach enables organizations to enforce unified security policies, optimize traffic routing, and ensure resilient connectivity.

Some of the SD-WAN capabilities supported on the Cisco Secure Firewall are:

- **Simplified management:**
  - SD-WAN Wizard
  - SASE: Umbrella auto tunnel deployment
  - Dynamic VTI (DVTI) hub spoke topology simplification
- **Application awareness:**
  - Direct Internet Access (DIA) for public cloud and guest user

- Policy based routing (PBR) using applications as a match criteria
- Local tunnel ID support for Umbrella

- **Increased usable bandwidth:**
  - ECMP support for load balancing across multiple ISPs and VTIs
  - Application-based load balancing using PBR

- **High availability with near zero network downtime:**
  - Dual ISP configuration
  - Optimal path selection based on application-based interface monitoring.

- **Secure Elastic Connectivity:**
  - Route-based (VTI) VPN tunnels between headquarters (hub) and branches (spokes)
  - IPv4 and IPv6 BGP, IPv4 and IPv6 OSPF, and IPv4 EIGRP over VTI
  - DVTI hubs that support spokes with static or dynamic IP

Features

The following table lists some commonly used SD-WAN features:

Feature	Introduced in	More Information
SD-WAN Wizard	Release 7.6	Using SD-WAN Wizard for Secure Branch Network Deployment, on page 2013
Application monitoring using SD-WAN Summary dashboard	Release 7.4.1	SD-WAN Summary Dashboard, on page 1705
SD-WAN Summary Dashboard	Release 7.4	SD-WAN Summary Dashboard, on page 1705
Policy-based routing with user identity and SGTs	Release 7.4	Policy Based Routing, on page 1323
Policy-based routing using HTTP path monitoring	Release 7.4	Policy Based Routing, on page 1323
Loopback interface support for VTIs	Release 7.3	About Loopback Interfaces, on page 829
Support for dynamic VTI (DVTI) with site-to-site VPN	Release 7.3	Dynamic VTI, on page 1531
Umbrella auto tunnel	Release 7.3	Deploy a SASE Tunnel on Umbrella, on page 1555

Feature	Introduced in	More Information
Support for IPv4 and IPv6 BGP, IPv4 and IPv6 OSPF, and IPv4 EIGRP for VTIs	Release 7.3	BGP, on page 1277 OSPF, on page 1239 EIGRP, on page 1267
Route-based site-to-site VPN with hub and spoke topology	Release 7.2	Create a Route-based Site-to-Site VPN, on page 1538
Policy-based routing with path monitoring	Release 7.2	Policy Based Routing, on page 1323
Site to Site VPN Monitoring Dashboard	Release 7.1	Monitoring the Site-to-Site VPNs, on page 1562
Direct Internet Access/Policy Based Routing	Release 7.1	Policy Based Routing, on page 1323
Equal-Cost-Multi-Path (ECMP) zone with WAN interfaces	Release 7.1	ECMP, on page 1223
ECMP zone with VTI interfaces	Release 7.1	ECMP, on page 1223
Backup VTI for route-based site-to-site VPN	Release 7.0	Route Traffic Through a Backup VTI Tunnel, on page 1549
Support for static VTI (SVTI) with site-to-site VPN	Release 6.7	Static VTI, on page 1530

Using SD-WAN Wizard for Secure Branch Network Deployment

Management Center allows you to easily configure VPN tunnels and routing configuration between your centralized headquarters (hubs) and remote branch sites (spokes) using the new SD-WAN wizard.

What are Hubs and Spokes?

Hubs: Devices that enable secure VPN connectivity to and from one or more remote branch devices or spokes. Hubs also act as a gateway for spokes to communicate with each other.

Spokes: Devices in remote branches that connect over VPN to a hub to securely access the corporate resources behind the hub. Spokes communicate with each other through the hub.

Benefits of Using SD-WAN Wizard

- Simplifies and automates the VPN and routing configuration of your SD-WAN network.
- Creates route-based VPN tunnels and simplifies the configuration process by automating tasks such as:
  - Generating tunnel interfaces of the branches.
  - Assigning IP addresses to the tunnel interfaces.

    - Configuring BGP for the SD-WAN overlay network. These configurations ensure seamless connection between hubs and spokes, and spoke to spoke through the hub.

- Provides seamless routing because hubs act as route reflectors and enable the following:
    - Provide connectivity between the spokes.
    - Determine the best routing path based on the spokes' active and backup tunnels.
- Requires minimal user input.
- Easily add multiple branches at a time.
- Provides easy dual ISP configurations.
- Enables network scaling.

Guidelines and Limitations for Using SD-WAN Wizard

Guidelines

- When you configure the DVTIs of two hubs, ensure that they have the same IPsec tunnel mode (IPv4 or IPv6).
- In a dual-hub SD-WAN topology, the hubs can be in different geographic locations and have different protected networks behind them. To ensure direct communication between these networks, ensure that you configure the following:
    - A point-to-point route-based VPN topology between the two hubs (**Devices** > **VPN** > **Site to Site**, and click **Add** > **Route-Based VPN**).
    - A dynamic routing protocol between the hubs (**Devices** > **Device Management**, click the device name and click **Routing**).
- When you configure IP address pools for spokes, ensure the following:
    - The **Allow Overrides** check box must be unchecked.
    - If you are using multiple pools, the IP addresses of the pools must not overlap.
    - IP addresses must not overlap with any of the interfaces on the spoke.
- When you create security zones or interface groups, choose **Routed** as the **Interface Type**.
- Use the spoke security zone to configure an access control policy that allows tunnel traffic to and from the spokes.
- Configure the spokes' static VTIs in an ECMP zone to load balance the application traffic. If you do not configure the ECMP zone, the remaining paths act as backup paths when the primary path goes down. Note that you must configure the spokes' static VTIs and not the physical interfaces in the ECMP zone. This configuration is not part of the SD-WAN wizard.
- In SD-WAN topologies with dual ISPs on spokes, the tunnel identity and the tunnel source of the spokes must be unique.

- If a spoke is part of multiple SD-WAN topologies, ensure that you use the same local community tag and learned route community tag in each SD-WAN topology. Note that the local community tags and learned route community tags must be different from each other.
- If a device has only IPv6 address configurations, you must configure the BGP router ID with a loopback or physical interface that has an IPv4 address (**Devices** > **Device Management**, and click **Routing** > **General Settings** > **BGP**).
- Configure unique local IKE identity for all tunnels across all your SD-WAN VPN topologies.
- Ensure that the spokes in an SD-WAN topology do not have the same protected network.
- Use distinct DVTIs for SD-WAN and route-based VPNs to avoid IPsec profile conflicts and errors.

Limitations

- You can configure a maximum of two hubs in an SD-WAN topology using the SD-WAN wizard.
- For each spoke, you can use only one WAN interface per topology. However, for dual-ISP setups, you can configure a second SD-WAN topology with the second WAN interface. For more information, see Sample Configurations for Dual ISP Deployment Using SD-WAN Wizard, on page 2020.
- SD-WAN wizard does not support the following:
  - IKEv1
  - Cluster devices are not supported on the hub and spoke because VTI is not supported on cluster devices.
  - Extranet hubs and spokes such as ASA, Cisco IOS, Cisco Viptela, Umbrella, Meraki, or vendor devices.

Prerequisites for Using the SD-WAN Wizard

- Firewall Management Center Essentials (formerly Base) license must allow export-controlled functionality.

  Choose **System > Licenses > Smart Licenses** to verify this functionality in the Firewall Management Center.
- You must be an Admin user.
- Hub devices must be Version 7.6.0 and later.
- Spoke devices must be Version 7.3.0 and later.
- The Firewall Threat Defense devices must have an internet-routable public IP address. The IP address can be static or dynamic.
- Assign appropriate logical names and IP addresses to the interfaces of the Firewall Threat Defense devices. For example, use *inside* for the interface connected to the LAN, and *outside* for the interface connected to the internet or WAN.
- If you are using certificate-based authentication, you must enroll the certificates in the hub and spokes.
- Configure routing, NAT, and AC policies to ensure underlay connectivity between the devices.

Configure an SD-WAN Topology Using the SD-WAN Wizard

The SD-WAN wizard allows you to easily configure VPN tunnels between your centralized headquarters and remote branch sites.

Before you begin

Ensure that you review Prerequisites for Using the SD-WAN Wizard, on page 2015 and Guidelines and Limitations for Using SD-WAN Wizard, on page 2014.

Procedure

Step 1 Choose **Devices** > **VPN** > **Site to Site**, and click **Add**.

Step 2 Enter a name for the SD-WAN VPN topology in the **Topology Name** field.

Step 3 Click the **SD-WAN Topology** radio button and click **Create**.

Step 4 Configure a hub:

a) Click **Add Hub**.

b) From the **Device** drop-down list, choose a hub.

c) Click **+** next to the **Dynamic Virtual Tunnel Interface (DVTI)** drop-down list to add a dynamic VTI for the hub.

The **Add Virtual Tunnel Interface** dialog box is prepopulated with default configurations. However, you must configure the **Tunnel Source**, and the **Borrow IP Address**. For more information, see Add a Dynamic Virtual Tunnel Interface for a Hub, on page 2019.

d) Click **OK**.

e) In the **Hub Gateway IP Address** field, enter the public IP address of the hub's VPN interface or the tunnel source of the dynamic VTI to which the spokes connect.

This IP address is auto populated if the interface has a static IP address. If hub is behind a NAT device, you must manually configure the post-NAT IP address.

f) From the **Spoke Tunnel IP Address Pool** drop-down list, choose an IP address pool or click **+** to create an address pool.

When you add spokes, the wizard auto generates spoke tunnel interfaces, and assigns IP addresses to these spoke interfaces from this IP address pool.

g) Click **Add** to save the hub configuration.

h) (Optional) To add a secondary hub, repeat Step 4a to Step 4f.

i) Click **Next**.

Step 5 Configure spokes:

Click **Add Spoke** to add a single spoke device, or click **Add Spokes (Bulk Addition)** to add multiple spokes to your topology.

- Click **Add Spoke**. In the **Add Spoke** dialog box, configure the following parameters:

  a. From the **Device** drop-down list, choose a spoke.

  b. From the **VPN Interface** drop-down list, choose a WAN-facing or internet-facing physical interface to establish a VPN connection with the hub.

c. Check the **Local Tunnel (IKE) Identity** check box to enable a unique and configurable identity for the VPN tunnel from this device to the remote peer. By default, this option is enabled.

d. Choose one of the following options from the **Identity Type** drop-down list:

- **Key ID**—(Default value) This value is auto populated as <*sd-wan topologyname*>_<*device_IP_address*>, for example, sdwantopo1_192.168.0.200. You can also specify a key ID of your choice.
- **Email ID**—Specify an email ID up to 127 characters.
- **IP Address**—IP address of the spoke's VPN interface.
- **Auto**—IP address of the spoke's VPN interface for pre-shared key authentication or the certificate Distinguished Name (DN) for certificate-based authentication.
- **Hostname**—Fully qualified hostname of the spoke.

e. Click **Save** to save the spoke configuration.

- Click **Add Spokes (Bulk Addition)**. In the **Add Bulk Spokes** dialog box, configure the following parameters:

a. Choose one or more devices from the **Available Devices** list and click **Add** to move the devices to **Selected Devices**.

b. Use one of the following methods to select the VPN interfaces of the spokes:

- Click the **Interface Name Pattern** radio button and specify a string to match the logical name of the internet or WAN interface of the spokes, for example, outside*, wan*.

  If the spoke has multiple interfaces with the same pattern, the first interface that matches the pattern is selected for the topology.

- Click the **Security Zone** radio button and choose a security zone with the VPN interfaces of the spokes from the drop-down list, or click + to create a security zone.

c. Click **Next**.

The wizard validates if the spokes have interfaces with the specified pattern. Only the validated devices are added to the topology.

d. Click **Add**.

e. Click **Next**.

For each spoke, the wizard automatically selects the hub's DVTI as the tunnel source IP address.

Note
If the hub’s tunnel source IP address is an IPv6 address, the wizard automatically selects the first IPv6 address of the spokes' selected interface. To edit the IPv6 address of a spoke's tunnel source, click the edit icon next to a spoke, choose an IPv6 address from the **IP Address** drop-down list, and click **Save**.

Step 6 Configure authentication settings for the devices in the SD-WAN topology:

a) **Authentication Type**—For device authentication, you can use a manual pre-shared key, an auto-generated pre-shared key, or a certificate.

- **Pre-shared Manual Key**—Specify the pre-shared key for the VPN connection.
- **Pre-shared Automatic Key**—(Default value) The wizard automatically defines the pre-shared key for the VPN connection. Specify the key length in the **Pre-shared Key Length** field. The range is 1 to 127.
- **Certificate**—When you use certificates as the authentication method, the peers obtain digital certificates from a CA server in your PKI infrastructure, and use them to authenticate each other.

b) Choose one or more algorithms from the **Transform Sets** drop-down list.
c) Choose one or more algorithms from the **IKEv2 Policies** drop-down list.
d) Click **Next**.

Step 7 Configure the SD-WAN settings:

This step involves the auto generation of spoke tunnel interfaces, and BGP configuration of the overlay network.

a) From the **Spoke Tunnel Interface Security Zone** drop-down list, choose a security zone or click **+** to create a security zone to which the wizard automatically adds the spokes' auto-generated Static Virtual Tunnel Interfaces (SVTIs).
b) Check the **Enable BGP on the VPN Overlay Topology** check box to automate BGP configurations such as neighbor configurations between the overlay tunnel interfaces and basic route redistribution from the directly connected LAN interfaces of the hubs and spokes.
c) In the **Autonomous System Number** field, enter an Autonomous System (AS) number.

AS number is a unique number for a network with a single routing policy. BGP uses AS numbers to identify networks. The spoke's BGP neighbor configuration is generated based on the corresponding hub's AS number. Range is from 0 to 65536.

- If all the hubs and spokes are in the same region, by default, **64512** is the AS number.
- If the primary and secondary hubs are in different regions, the primary hub and the spokes are configured with **64512** as the AS number, and the secondary hub is configured with a different AS number.

d) In the **Community Tag for Local Routes** field, enter the BGP community attribute to tag connected and redistributed local routes. This attribute enables easy route filtering.
e) Check the **Redistribute Connected Interfaces** check box and choose an interface group from the drop-down list or click **+** to create an interface group with connected inside or LAN interfaces for BGP route redistribution in the overlay topology.
f) Check the **Enable Multiple Paths for BGP** check box to allow multiple BGP routes to be used at the same time to reach the same destination. This option enables BGP to load-balance traffic across multiple links.

Note that when you enable this option, BGP multipath is enabled only for spokes.

g) (Optional) Check the **Secondary Hub is in Different Autonomous System** check box. This check box appears only if you have a secondary hub in this topology.
h) In the **Autonomous System Number** field, enter the AS number for the secondary hub.
i) In the **Community Tag for Learned Routes** field, enter the BGP community attribute to tag routes learned from other SD-WAN peers over the VPN tunnel. This attribute is required only for eBGP configuration when the secondary hub has a different AS number. This field appears only if you have configured two hubs in the SD-WAN topology.

j) Click **Next**.

Step 8 Click **Finish** to save and validate the SD-WAN topology.

You can view the topology in the **Site-to-Site VPN Summary** page (**Devices** > **VPN** > **Site to Site**). After you deploy the configurations to all the devices, you can see the status of all the tunnels in this page.

What to do next

- View the auto-generated spoke SVTIs and their IP addresses—Click the edit icon next to the spoke configuration and click **View Generated Tunnel Interfaces**.
- We recommend that you enable ECMP on the spoke SVTIs. Choose **Devices** > **Device Management**, and click **Routing** > **ECMP**.
- Deploy the configurations on the hub and spokes. Choose **Deploy**. Select devices and click **Deploy**.
- Verify the SD-WAN topology tunnel statuses. For more information, see Verify Tunnel Statuses of an SD-WAN Topology, on page 2025.
- Configure ACLs for the spokes' tunnel interface security zones. Choose **Policies > Access Control**.
- We recommend that you enable BGP multipath on hubs. To enable BGP multipath on hubs:
  1. Choose ECMP**Devices** > **Device Management**, and click **Routing**.
  2. Under **General Settings**, click **BGP**.
  3. Check the **Enable BGP** check box to enable BGP.
  4. In the **AS Number** field, enter the AS number that you configured in the SD-WAN topology.
  5. Click **Save**.
  6. In the left pane, choose **BGP > IPv4 or IPv6**, and click the **General** tab.
  7. In the **Forward Packets over Multiple Paths** section, click the edit icon.
  8. Configure values for **Number of Paths** and **IBGP number of paths**. We recommend that you configure these values as 8.
- For more information about configuration examples using SD-WAN wizard, see Sample Configurations for Dual ISP Deployment Using SD-WAN Wizard, on page 2020
- Configure a PBR policy on each spoke for application-aware routing based on the application performance metrics of the WAN interfaces. For more information, see Route Application Traffic from the Branch to the Internet Using Direct Internet Access (DIA).

Add a Dynamic Virtual Tunnel Interface for a Hub

In the SD-WAN wizard, you must configure a DVTI for each hub. DVTI uses a virtual template to dynamically generate a unique virtual access interface for each VPN session.

Before you begin

In the SD-WAN wizard, click **Add Hub**, and choose a hub from the **Device** drop-down list.

Procedure

Step 1 Click + next to the **Dynamic Virtual Tunnel Interface (DVTI)** drop-down list to add a DVTI for the hub.

The **Add Virtual Tunnel Interface** dialog box appears with the following prepopulated default configurations.

a. **Tunnel Type**: Dynamic.

b. **Name**: <*tunnel_source_interface_logical name*>_dynamic_vti_<*tunnel_ID*>. For example, outside_dynamic_vti_1.

c. **Enabled** check box: Checked by default.

d. **Template ID**: Unique ID for the DVTI.

e. **Tunnel Source**: Physical interface that is the source of the DVTI and is auto populated by default.

f. **IPsec Tunnel Mode**: IPv4, by default.

Step 2 Choose a security zone for the dynamic VTI from the **Security Zone** drop-down list.

Step 3 Choose a physical or loopback interface from the **Borrow IP** drop-down list, the dynamic VTI interface inherits this IP address.

Ensure that you use an IP address different from the tunnel source IP address. We recommend that you use a loopback IP address.

Step 4 Click **OK** to save the dynamic VTI.

Sample Configurations for Dual ISP Deployment Using SD-WAN Wizard

Dual ISP Deployment: Two Hubs and Four Spokes in the Same Region

In the following dual ISP topology, the hubs and the spokes are in a single region, with AS number as 1111. The hubs and spokes use Internal Border Gateway Protocol (iBGP) as the routing protocol to exchange routing information.

- Hub HA1 and Hub HA2 are hub threat defense devices at the headquarters.
- Branch1, Branch2, Branch3, and Branch4 are spoke threat defense devices at the branches.
- ISP1 is the VPN interface of each spoke to ISP1.
- ISP2 is the VPN interface of each spoke to ISP2.

Figure 444: Dual ISP Topology with Two Hubs and Four Spokes in the Same Region

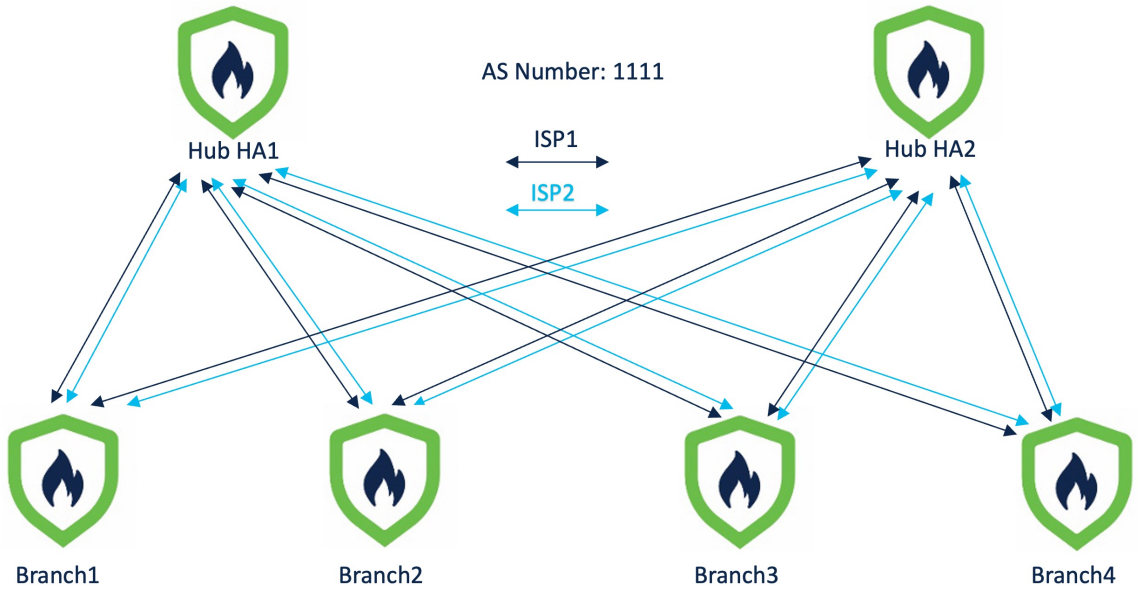

To configure this topology, you must create the following two SD-WAN topologies using the SD-WAN wizard:

SD-WAN Topology 1

Parameter	Value
Primary Hub	Hub HA1
Secondary Hub	Hub HA2
Spokes	Branch1, Branch2, Branch3, Branch4
AS Number	1111
VPN Interface (Spoke Tunnel Source)	ISP1
Number of Tunnels	8

The total number of tunnels in SD-WAN Topology 1 is 8.

SD-WAN Topology 2

Parameter	Value
Primary Hub	Hub HA1
Secondary Hub	Hub HA2
Spokes	Branch1, Branch2, Branch3, Branch4
AS Number	1111

Parameter	Value
VPN Interface (Spoke Tunnel Source)	ISP2
Number of Tunnels	8

The total number of tunnels in SD-WAN Topology 2 is 8.

The total number of VPN tunnels for this dual ISP deployment is 16.

Note If the hubs are in different geographic locations and have different protected networks behind them, to ensure direct communication between these networks, configure a point-to-point route-based VPN topology between the two hubs using the route-based VPN wizard.

Dual ISP Deployment: Two Hubs and Four Spokes in Different Regions

In the following dual ISP topology, the hubs are in different regions, and have two directly connected spokes each. The hubs and their directly connected spokes use Internal Border Gateway Protocol (iBGP) as the routing protocol, and the hubs use External Border Gateway Protocol (eBGP) to exchange routing information.

- Hub HA1 and Hub HA2 are hub threat defense devices at the headquarters.
- Branch1, Branch2, Branch3, and Branch4 are spoke threat defense devices at the branches.
- HQ1, Branch1, and Branch2 are in a single region with AS number as 1111.
- HQ2, Branch3, and Branch4 are in a single region with AS number as 2222.
- ISP1 is the VPN interface of each spoke to ISP1.
- ISP2 is the VPN interface of each spoke to ISP2.

Figure 445: Dual ISP Topology with Two Hubs and Four Spokes in Different Regions

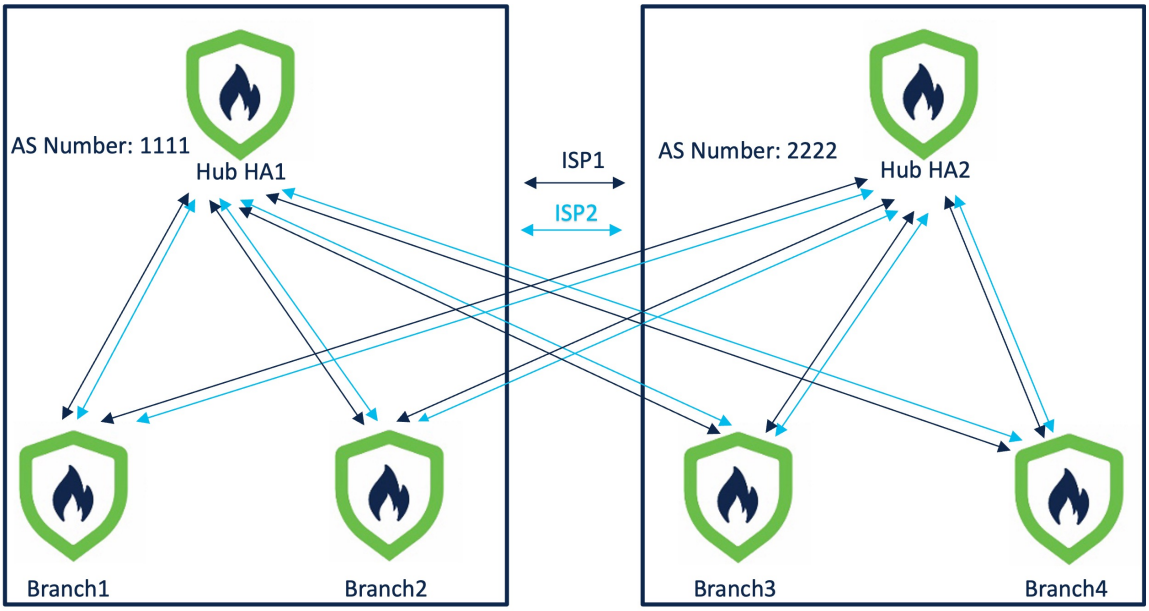

To configure this topology, you must create the following four SD-WAN topologies using the SD-WAN wizard:

SD-WAN Topology 1

Parameter	Value
Primary Hub	Hub HA1
Secondary Hub	Hub HA2
Spokes	Branch1, Branch2
AS Number	1111
Secondary AS Number	2222
VPN Interface (Spoke Tunnel Source)	ISP1

The number of tunnels in SD-WAN Topology 1 is 4.

SD-WAN Topology 2

Parameter	Value
Primary Hub	Hub HA1
Secondary Hub	Hub HA2
Spokes	Branch1, Branch2
AS Number	1111

Parameter	Value
Secondary AS Number	2222
VPN Interface (Spoke Tunnel Source)	ISP2

The number of tunnels in SD-WAN Topology 2 is 4.

SD-WAN Topology 3

Parameter	Value
Primary Hub	Hub HA2
Secondary Hub	Hub HA1
Spokes	Branch3, Branch4
AS Number	2222
Secondary AS Number	1111
VPN Interface (Spoke Tunnel Source)	ISP1

The number of tunnels in SD-WAN Topology 3 is 4.

SD-WAN Topology 4

Parameter	Value
Primary Hub	Hub HA2
Secondary Hub	Hub HA1
Spokes	Branch3, Branch4
AS Number	2222
Secondary AS Number	1111
VPN Interface (Spoke Tunnel Source)	ISP2

The number of tunnels in SD-WAN Topology 4 is 4.

The total number of VPN tunnels for this dual ISP deployment is 16.

Note If the hubs are in different geographic locations and have different protected networks behind them, to ensure direct communication between these networks, configure a point-to-point route-based VPN topology between the two hubs using the route-based VPN wizard.

Verify Tunnel Statuses of an SD-WAN Topology

Verify Tunnel Statuses on the Site-to-Site VPN Summary Page

To verify if the VPN tunnels of the SD-WAN topologies are up, choose **Devices** > **VPN** > **Site to Site**.

Following are the five SD-WAN topologies with two hubs and four spokes in different regions that are connected to dual ISPs:

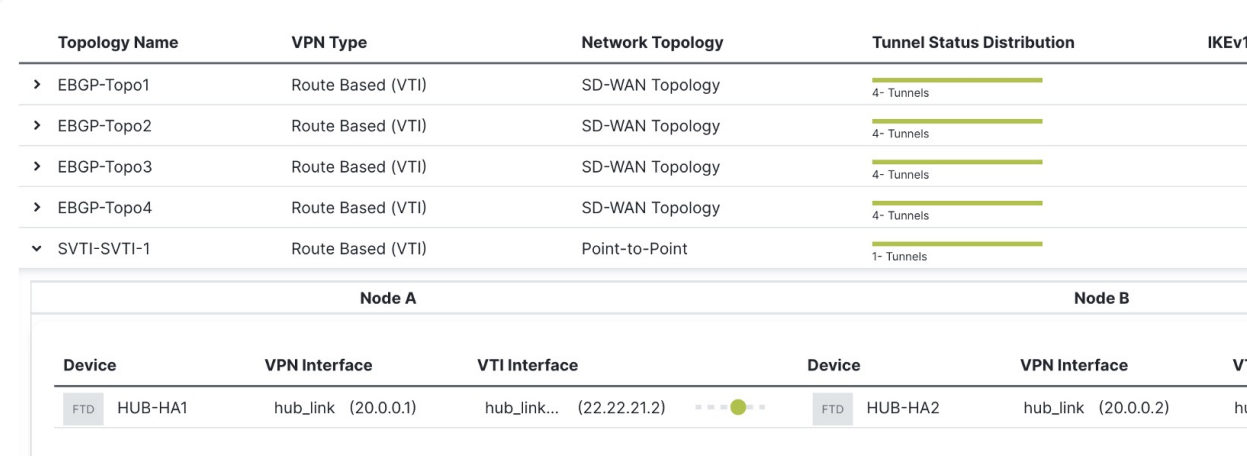

Verify Tunnel Statuses on the Site-to-Site VPN Dashboard

To view details of the SD-WAN VPN tunnels, choose **Overview** > **Dashboards** > **Site to Site VPN**.

Following are the VPN tunnels for an SD-WAN topology with two hubs and four spokes in different regions that are connected to dual ISPs:

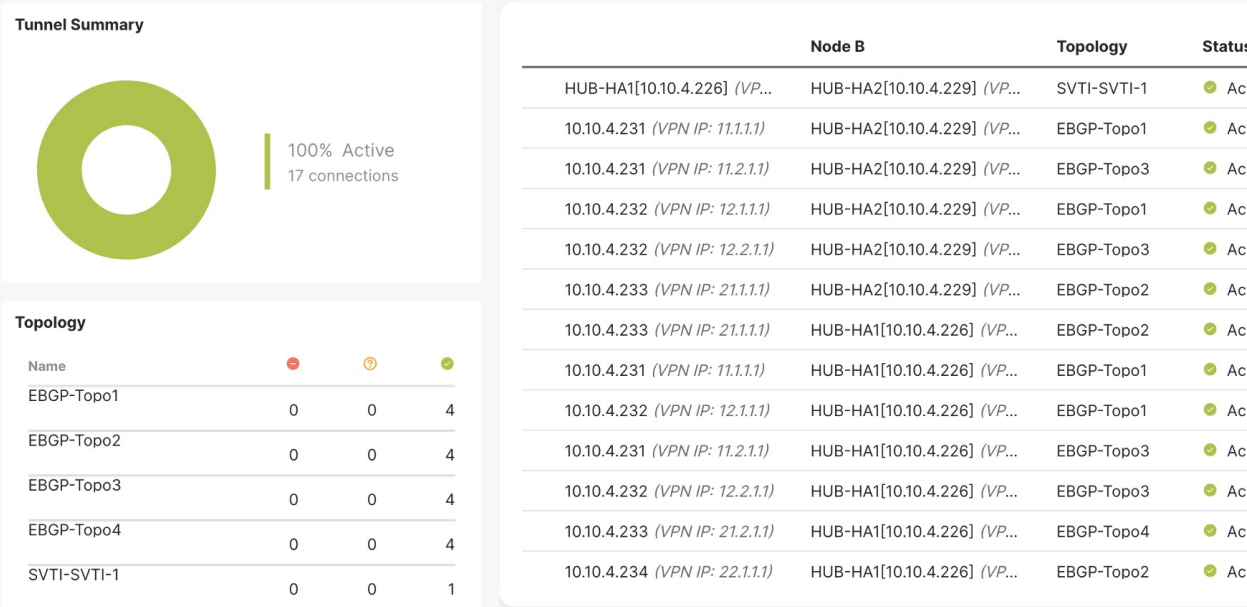

To view more details of each VPN tunnel:

1. Hover over a tunnel.
2. Click the **View Full Information** ( ) icon. A pane with tunnel details and more actions appears.
3. Click the **CLI Details** tab in the side pane to view the show commands and details of the IPsec security associations.

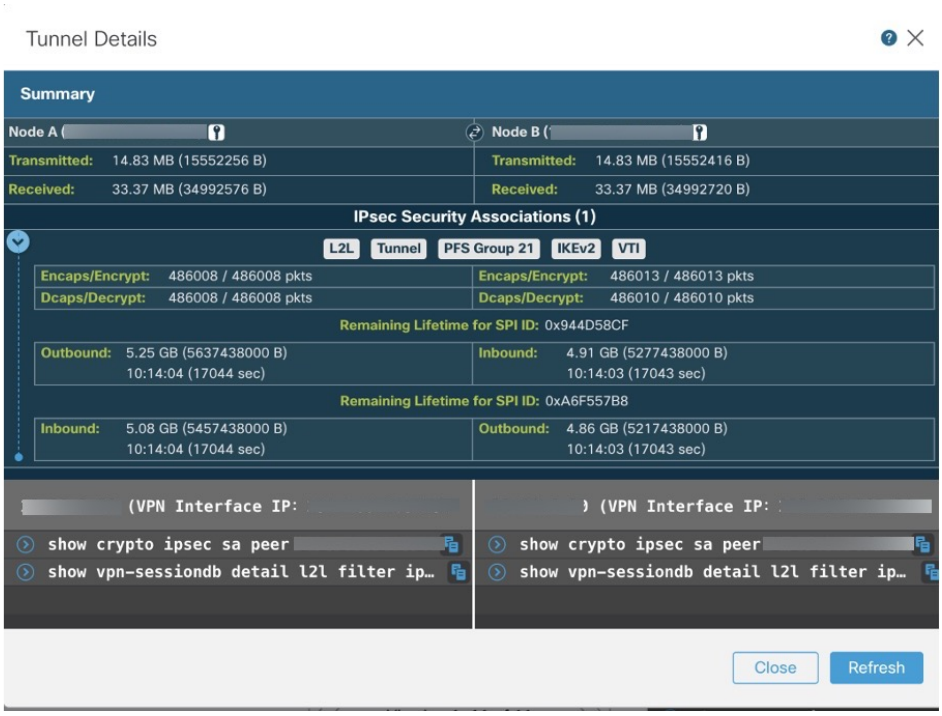

View Virtual Tunnel Interfaces of the Devices

To view the dynamic VTIs of hubs and static VTIs of spokes:

1. Choose **Devices** > **Device Management**.
2. Click the edit icon for a hub or a spoke device.
3. Click the **Interface** tab.
4. Click the **Virtual Tunnels** tab.

   For each VTI, you can view details such as name, IP address, IPsec mode, tunnel source interface details, topology, and remote peer IP.

Following image shows an example of the virtual access interfaces created dynamically by a hub's DVTI:

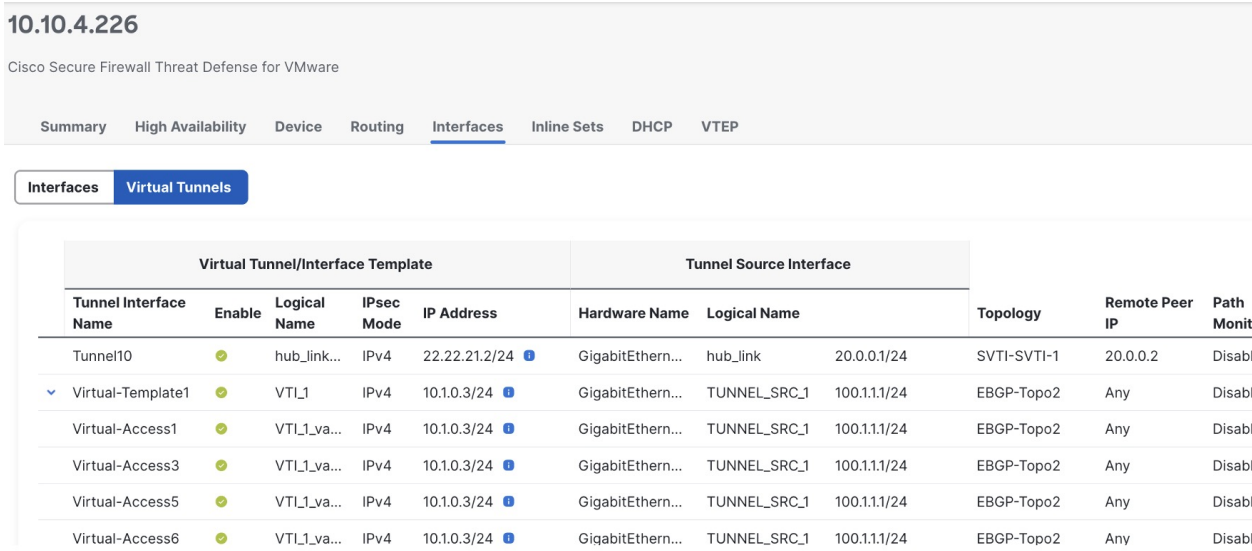

Following image shows an example of the static tunnel virtual interfaces (SVTIs) created on a spoke by the SD-WAN wizard:

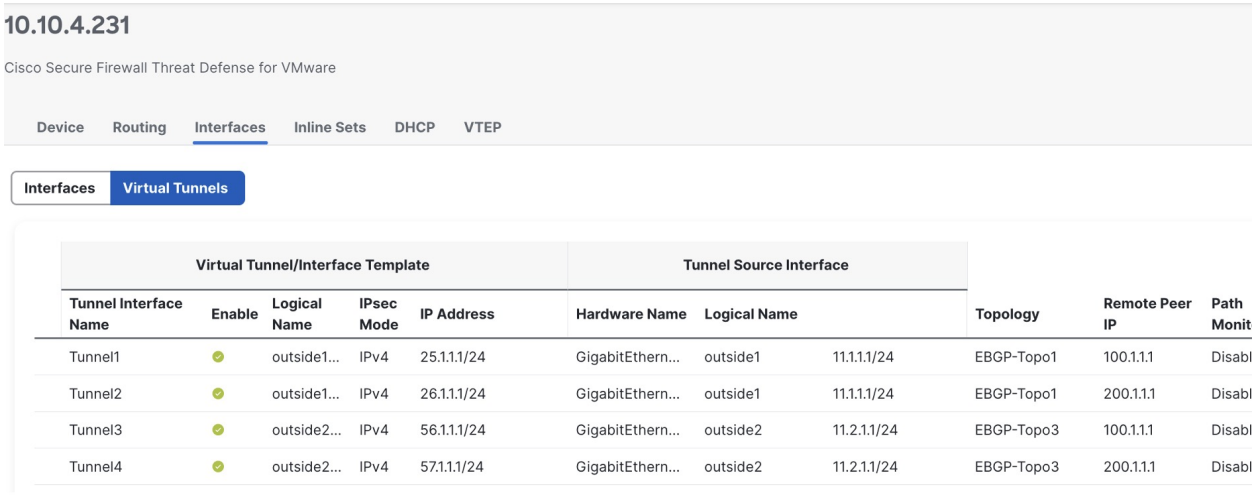

The SD-WAN wizard assigns IP addresses to these tunnel interfaces from the IP address pool of the hub.

Verify Routing on the Hub and Branches

To verify the BGP configuration of the hubs and spokes of the SD-WAN topologies:

1. Choose **Devices** > **Device Management**.
2. Click the edit icon for a hub or a spoke device.
3. Click the **Device** tab.
4. Click **CLI** in the **General** card. The **CLI Troubleshoot** window appears.
5. Enter the following commands in the **Command** field and click **Execute**:
   - **show route**

- **show bgp summary**

Use Cases for SD-WAN Capabilities

- Simplify Branch to Hub Communication using Dynamic Virtual Tunnel Interface (DVTI)
- Route Application Traffic from the Branch to the Internet Using Direct Internet Access (DIA)
- Secure Internet Traffic Using Umbrella Auto Tunnel

PART IX

Intrusion Detection and Prevention

- An Overview of Network Analysis and Intrusion Policies, on page 2031
- Migrate from Snort 2 to Snort 3, on page 2049
- Get Started with Snort 3 Intrusion Policies, on page 2059
- Tune Intrusion Policies Using Rules, on page 2073
- Tailor Intrusion Protection for Your Network Assets, on page 2093
- Intrusion Prevention Performance Tuning, on page 2099
- Get Started with Snort 3 Network Analysis Policies, on page 2113
- Encrypted Visibility Engine, on page 2147
- Elephant Flow Detection, on page 2155
- Use Case - Migrate from Snort 2 to Snort 3 In Secure Firewall Management Center, on page 2159
- Use Case - Generate Snort 3 Recommendations In Secure Firewall Management Center, on page 2171
- Use Case - Block Traffic Based on the EVE Threat Confidence Score, on page 2179
- Use Case - Configure Elephant Flow Detection Outcomes, on page 2187
- Mitigate Threats Using MITRE Framework in Snort 3 Intrusion Policies, on page 2197

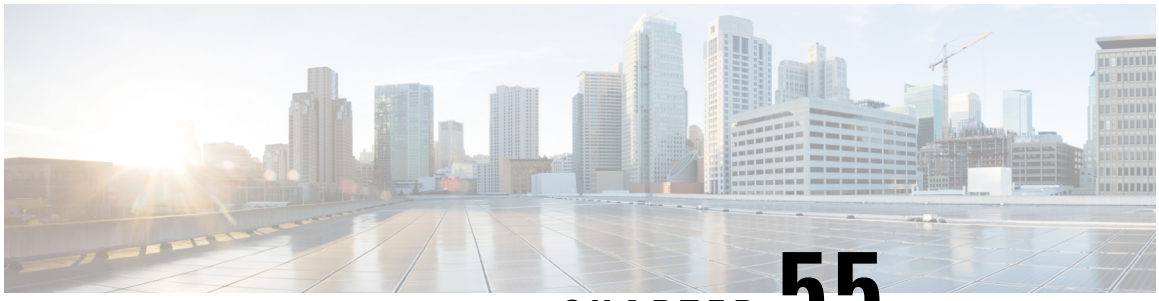

CHAPTER 55

An Overview of Network Analysis and Intrusion Policies

The Snort inspection engine is an integral part of the Secure Firewall Threat Defense (formerly Firepower Threat Defense) device. This chapter provides an overview of Snort 3 and the network analysis and intrusion policies. It also provides an insight into system-provided and custom network analysis and intrusion policies.

- About Network Analysis and Intrusion Policies, on page 2031
- Snort Inspection Engine, on page 2032
- Snort 3, on page 2033
- Guidelines and Limitations for Network Analysis and Intrusion Policies, on page 2035
- How Policies Examine Traffic For Intrusions, on page 2036
- System-Provided and Custom Network Analysis and Intrusion Policies, on page 2042
- Prerequisites for Network Analysis and Intrusion Policies, on page 2048

About Network Analysis and Intrusion Policies

Network analysis and intrusion policies work together as part of the intrusion detection and prevention feature.

- The term *intrusion detection* generally refers to the process of passively monitoring and analyzing network traffic for potential intrusions and storing attack data for security analysis. This is sometimes referred to as "IDS."
- The term *intrusion prevention* includes the concept of intrusion detection, but adds the ability to block or alter malicious traffic as it travels across your network. This is sometimes referred to as "IPS."

In an intrusion prevention deployment, when the system examines packets:

- A **network analysis policy** governs how traffic is *decoded* and *preprocessed* so it can be further evaluated, especially for anomalous traffic that might signal an intrusion attempt.
- An **intrusion policy** uses *intrusion and preprocessor rules* (sometimes referred to collectively as *intrusion rules*) to examine the decoded packets for attacks based on patterns. Intrusion policies are paired with *variable sets*, which allow you to use named values to accurately reflect your network environment.

Both network analysis and intrusion policies are invoked by a parent access control policy, but at different times. As the system analyzes traffic, the network analysis (decoding and preprocessing) phase occurs before and separately from the intrusion prevention (additional preprocessing and intrusion rules) phase. Together,

network analysis and intrusion policies provide broad and deep packet inspection. They can help you detect, alert on, and protect against network traffic that could threaten the availability, integrity, and confidentiality of hosts and their data.

The system is delivered with several similarly named network analysis and intrusion policies (for example, Balanced Security and Connectivity) that complement and work with each other. By using system-provided policies, you can take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos sets intrusion and inspector rule states, as well as provides the initial configurations for inspectors and other advanced settings.

You can also create custom network analysis and intrusion policies. You can tune settings in custom policies to inspect traffic in the way that matters most to you so that you can improve both the performance of your managed devices and your ability to respond effectively to the events they generate.

You create, edit, save, and manage network analysis and intrusion policies using similar policy editors in the web interface. When you are editing either type of policy, a navigation panel appears on the left side of the web interface; the right side displays various configuration pages.

Refer to the videos for additional support and information:

- Snort 3 Condensed Overview
- Snort 3 Extended Overview

Attention **Detection mode deprecation:** From management center Version 7.4.0, for a network analysis policy (NAP), the **Detection** inspection mode is deprecated and will be removed in an upcoming release.

The **Detection** mode was intended to be used as a test mode so that you can enable inspections and see how they behave in your network before setting it to drop traffic, that is, to show traffic that would be dropped.

This behavior is improved where all inspector drops are controlled by the rule state, and you can set each one to generate events. This is done to test them before configuring the rule state to drop traffic. As we now have granular control over traffic drops in Snort 3, the **Detection** mode only adds more complexity to the product and is not needed, so the detection mode is deprecated.

If you change a NAP in **Detection** mode to **Prevention**, the NAP that processes the traffic of intrusion events and have the result "will be dropped" will now be "dropped" and the corresponding traffic will drop the traffic from these events. This is applicable for rules whose GIDs are not 1 or 3. GIDs 1 and 3 are text/compiled rules (typically provided by Talos or from your custom/imported rules) and all other GIDs are inspections for anomalies. These are more uncommon rules to trigger in a network. Changing to **Prevention** mode is unlikely to have any impact on the traffic. You need to just disable the intrusion rule that is applicable for the dropped traffic and set it to just generate or disable.

We recommend you choose **Prevention** as the inspection mode, but if you choose **Prevention**, you cannot revert to **Detection** mode.

Snort Inspection Engine

The Snort inspection engine is an integral part of the Secure Firewall Threat Defense (formerly Firepower Threat Defense) device. The inspection engine analyzes traffic in real time to provide deep packet inspection. Network analysis and intrusion policies together utilize the Snort inspection engine's capabilities to detect and protect against intrusions.

Snort 3

Snort 3 is the latest version of the Snort inspection engine, which has vast improvements compared to the earlier version of Snort. The older version of Snort is Snort 2. Snort 3 is more efficient, and it provides better performance and scalability.

Note Snort 2 is not supported on threat defense Version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the Firewall Management Center guide that matches your Firewall Threat Defense version.

Snort 3 is architecturally redesigned to inspect more traffic with equivalent resources when compared to Snort 2. Snort 3 provides simplified and flexible insertion of traffic parsers. Snort 3 also provides new rule syntax that makes rule writing easier and shared object rule equivalents visible.

The other significant changes with Snort 3 are:

- Unlike Snort 2, which uses multiple Snort instances, Snort 3 associates multiple threads with a single Snort instance. This uses less memory, improves Snort reload times, and supports more intrusion rules and a larger network map. The number of Snort threads varies by platform and is the same as the number of Snort 2 instances for each platform. Usage is virtually transparent.
- Snort version per Firewall Threat Defense—The Snort inspection engine is Firewall Threat Defense specific and not Secure Firewall Management Center (formerly Firepower Management Center) specific. Firewall Management Center can manage several Firewall Threat Defenses, each with either versions of Snort (Snort 2 and Snort 3). Although the Firewall Management Center's intrusion policies are unique, the system applies Snort 2 or Snort 3 version of an intrusion policy for intrusion protection depending on the device's selected inspection engine.
- Decoder rules—Packet decoder rules fire only in the default intrusion policy. The system ignores decoder rules that you enable in other policies.
- Shared object rules—Snort 3 supports some but not all shared object (SO) intrusion rules (rules with a generator ID (GID) of 3). Enabled shared object rules that are not supported do not trigger.
- Multi-layer inspection for Security Intelligence—Snort 3 detects the innermost IP address regardless of the layer.
- Platform support—Snort 3 requires Firewall Threat Defense 7.0 or later. It is not supported with ASA FirePOWER or NGIPSv.
- Managed Devices—An Firewall Management Center with version 7.0 can simultaneously support version 6.4, 6.5, 6.6, 6.7, and 7.0 Snort 2 Firewall Threat Defenses, and version 7.0 Snort 3 Firewall Threat Defenses.
- Traffic interruption when switching Snort versions—Switching Snort versions interrupts traffic inspection and a few packets might drop during deployment.
- Unified policies—Irrespective of the underlying Snort engine version that is enabled in the managed Firewall Threat Defenses, the access control policies, intrusion policies, and network analysis policies configured in the Firewall Management Center work seamlessly in applying the policies. All intrusion policies in Firewall Management Center version 7.0 and above have two versions available, Snort 2 version and Snort 3 version. The intrusion policy is unified, which means that it has a common name,

base policy, and inspection mode, although there are two versions of the policy (Snort 2 version and Snort 3 version). The Snort 2 and the Snort 3 versions of the intrusion policy can be different in terms of the rule settings. However, when the intrusion policy is applied on a device, the system automatically identifies the Snort version enabled on the device and applies the rule settings configured for that version.

- Lightweight Security Package (LSP)—Replaces the Snort Rule Updates (SRU) for Snort 3 next-generation intrusion rule and configuration updates. Downloading updates downloads both the Snort 3 LSP and the Snort 2 SRU.

  LSP updates provide new and updated intrusion rules and inspector rules, modified states for existing rules, and modified default intrusion policy settings for Firewall Management Center and Firewall Threat Defense versions 7.0 or above. When you upgrade an Firewall Management Center from version 6.7 or lower to 7.0, it supports both LSPs and SRUs. LSP updates may also delete system-provided rules, provide new rule categories and default variables, and modify default variable values. For more information on LSP updates, see the *Update Intrusion Rules* topic in the latest version of the *Firepower Management Center Configuration Guide*.

- Mapping of Snort 2 and Snort 3 rules and presets—Snort 2 and Snort 3 rules are mapped and the mapping is system-provided. However, it is not a one-to-one mapping. The system-provided intrusion base policies are pre-configured for both Snort 2 and Snort 3, and they provide the same intrusion prevention although with different rule sets. The system-provided base policies for Snort 2 and Snort 3 are mapped with each other for the same intrusion prevention settings. For more information, see View Snort 2 and Snort 3 Base Policy Mapping, on page 2054.

- Synchronizing Snort 2 and Snort 3 rule override—When a Firewall Threat Defense is upgraded to 7.0, you can upgrade the inspection engine of the Firewall Threat Defense to the Snort 3 version. Firewall Management Center maps all the overrides in the existing rules of the Snort 2 version of the intrusion policies to the corresponding Snort 3 rules using the mapping provided by Talos. However, if there are additional overrides performed after the upgrade or if you have installed a new Firewall Threat Defense of version 7.0, they have to be manually synchronized. For more information, see Synchronize Snort 2 Rules with Snort 3 , on page 2054.

- Custom intrusion rules—You can create custom intrusion rules in Snort 3. You can also import the custom intrusion rules that exist for Snort 2 to Snort 3. For more information, see Custom Rules in Snort 3, on page 2075.

- Rule groups—The Firewall Management Center groups all Snort 3 rules into rule groups. Rule groups are logical groups of rules which provide an easy management interface to enhance rule accessibility, rule navigation, and a better control over the rule group security level.

  From Firewall Management Center 7.3.0, rule navigation for several levels of rule groups is supported that provides more flexibility and logical grouping of rules. The MITRE framework is added that enables you to navigate through rules using the MITRE framework. MITRE is just another category of rule groups and are a part of Talos rule groups.

Note See https://attack.mitre.org for information about MITRE.

A rule can be part of multiple rule groups, such as multiple MITRE ATT&CK rule groups, a rule category rule group, multiple "asset type" rule groups, a malware campaign, and others. The available rule groups are listed in the intrusion policy editor and can be selected to enhance policies.

With this multi-level hierarchical structure, you can traverse down to the last element, which is the “leaf rule group.” These rule groups contain sets of rules that are related to each other, such as a specific type

of vulnerability, a similar target system, or a similar threat category. Rule groups have four security levels associated with them. You can change the security level, add or remove rule groups, and you can change the rule action for rules that match traffic seen on the network. This is done to bring a satisfactory balance between security, performance, and false positive resistance.

To edit a Snort 3 intrusion policy, see Edit Snort 3 Intrusion Policies, on page 2061.

For rule group reporting in intrusion events, see Rule Group Reporting, on page 2065.

- Switching between Snort 2 and Snort 3 engines—Firewall Threat Defenses (version 7.6 and earlier) that support Snort 3 can also support Snort 2. Switching from Snort 3 to Snort 2 is not recommended from the efficacy perspective.

Important Although you can switch Snort versions freely, intrusion rule changes in one version of Snort will not be updated in the other version automatically. If you change the rule action for a rule in one version of Snort, ensure you replicate the change in the other version before switching the Snort version. System provided synchronization option only synchronizes the changes in the Snort 2 version of the intrusion policy to the Snort 3 version, and not the other way around.

Guidelines and Limitations for Network Analysis and Intrusion Policies

- A high percentage of traffic with small packets causes Snort performance to decrease. This behaviour is observed even when all the preprocessors are disabled.
- When you attempt to deploy a configuration change on a threat defense device with low memory, snort deployment is also triggered. This results in high consumption of RSS memory. Snort memory usage is also impacted if you deploy large configurations on the device, such as multiple IPS policies containing a large number of snort IPS rules, network objects, and access-control lists. You can mitigate such memory issues by optimizing the configuration. For best practices on how to configure access control rules to optimize the configuration, see Best Practices for Access Control Rules.
- If you increase the memory of a Threat Defense Virtual instance, you must redeploy the configuration for Snort 3 to utilize the additional memory.

Note The Snort 3 memory allocation is not automatically adjusted when you increase the memory of the Threat Defense Virtual instance. You must redeploy the configuration to regenerate relevant configuration files, such as `memory_allocation.lua`, which apply the updated resource limits to Snort 3.

- If an SIP stream is followed by RTP streams from the same connection, Snort inspects the initial SIP communication that is sent for connection establishment and allows SIP traffic. The RTP streams that follow the SIP communication are also trusted by default and bypass the configured rules. To prevent such scenarios, trusting the parent SIP connection or adding the parent SIP connection to a prefilter rule

ensures that only the SIP stream bypasses Snort inspection and allows the subsequent RTP streams to be evaluated separately against the corresponding rules.

- When intrusion packet events are forwarded via syslog or the eStreamer fully-qualified event feed, the packet data field may be truncated due to limitations in the buffer size available for syslog generation and the eStreamer. In such scenarios, the packet length will not match the actual packet data that is sent.
- You cannot make policy changes, switch snort versions, and deploy both these changes at the same time. You must make the required policy changes and deploy, or switch the snort versions and deploy.

Feature Limitations of Snort 3 for Firewall Management Center-Managed Firewall Threat Defense

The following table lists the features that are supported on Snort 2 but not supported on Snort 3 for Firewall Management Center-managed Firewall Threat Defense devices.

Table 108: Feature Limitations of Snort 3

Policy/Area	Features not supported
Access Control Policy	The following application settings: • Safe Search • YouTube EDU
Intrusion Policy	• Global rule thresholding • Logging configuration: • SNMP • SRU rule updates as Snort 3 supports only LSP rule updates
Other features	Event logging with FQDN names

How Policies Examine Traffic For Intrusions

When the system analyzes traffic as part of your access control deployment, the network analysis (decoding and preprocessing) phase occurs before and separately from the intrusion prevention (intrusion rules and advanced settings) phase.

The following diagram shows, in a simplified fashion, the order of traffic analysis in an inline, intrusion prevention and AMP for Networks deployment. It illustrates how the access control policy invokes other policies to examine traffic, and in which order those policies are invoked. The network analysis and intrusion policy selection phases are highlighted.

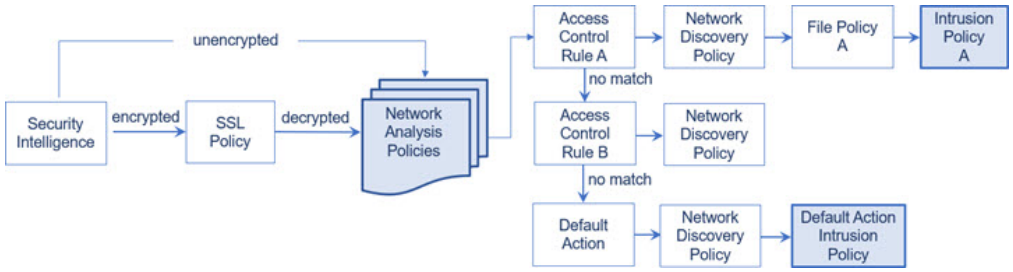

In an inline deployment (that is, where relevant configurations are deployed to devices using routed, switched, or transparent interfaces, or inline interface pairs), the system can block traffic without further inspection at almost any step in the illustrated process. Security Intelligence, the SSL policy, network analysis policies, file policies, and intrusion policies can all either drop or modify traffic. Only the network discovery policy, which passively inspects packets, cannot affect the flow of traffic.

Similarly, at each step of the process, a packet could cause the system to generate an event. Intrusion and preprocessor events (sometimes referred to collectively as *intrusion events*) are indications that a packet or its contents may represent a security risk.

Tip The diagram does not reflect that access control rules handle encrypted traffic when your SSL inspection configuration allows it to pass, or if you do not configure SSL inspection. By default, the system disables intrusion and file inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion and file inspection configured.

Note that for a single connection, although the system selects a network analysis policy before an access control rule as shown in the diagram, some preprocessing (notably application layer preprocessing) occurs after access control rule selection. This does **not** affect how you configure preprocessing in custom network analysis policies.

Decoding, Normalizing, and Preprocessing: Network Analysis Policies

Without decoding and preprocessing, the system could not appropriately evaluate traffic for intrusions because protocol differences would make pattern matching impossible. Network analysis policies govern these traffic-handling tasks:

- **after** traffic is filtered by Security Intelligence
- **after** encrypted traffic is decrypted by an optional SSL policy
- **before** traffic can be inspected by file or intrusion policies

A network analysis policy governs packet processing in phases. First the system decodes packets through the first three TCP/IP layers, then continues with normalizing, preprocessing, and detecting protocol anomalies:

- The packet decoder converts packet headers and payloads into a format that can be easily used by the inspectors and later, intrusion rules. Each layer of the TCP/IP stack is decoded in turn, beginning with the data link layer and continuing through the network and transport layers. The packet decoder also detects various anomalous behaviors in packet headers.
- In inline deployments, the inline normalization preprocessor reformats (normalizes) traffic to minimize the chances of attackers evading detection. It prepares packets for examination by other inspectors and

intrusion rules, and helps ensure that the packets the system processes are the same as the packets received by the hosts on your network.

- Various network and transport layers inspectors detect attacks that exploit IP fragmentation, perform checksum validation, and perform TCP and UDP session preprocessing.

  Note that some advanced transport and network inspector settings apply globally to all traffic handled by the target devices of an access control policy. You configure these in the access control policy rather than in a network analysis policy.

- Various application-layer protocol decoders normalize specific types of packet data into formats that the intrusion rules engine can analyze. Normalizing application-layer protocol encodings allows the system to effectively apply the same content-related intrusion rules to packets whose data is represented differently, and to obtain meaningful results.

- The Modbus, DNP3, CIP, and s7commplus SCADA inspectors detect traffic anomalies and provide data to intrusion rules. Supervisory Control and Data Acquisition (SCADA) protocols monitor, control, and acquire data from industrial, infrastructure, and facility processes such as manufacturing, production, water treatment, electric power distribution, airport and shipping systems, and so on.

- Several inspectors allow you to detect specific threats, such as Back Orifice, portscans, SYN floods and other rate-based attacks.

  Note that you configure the sensitive data inspector, which detects sensitive data such as credit card numbers and Social Security numbers in ASCII text, in intrusion policies.

Note When TLS Server Identity is disabled, Snort 3 does not perform SNI mismatch detection. It only evaluates the SNI in the Client Hello packet and bypasses the validation of the certificate's Common Name (CN) in the Server Hello packet.

In a newly created access control policy, one default network analysis policy governs preprocessing for *all* traffic for *all* intrusion policies invoked by the same parent access control policy. Initially, the system uses the Balanced Security and Connectivity network analysis policy as the default, but you can change it to another system-provided or custom network analysis policy. In a more complex deployment, advanced users can tailor traffic preprocessing options to specific security zones, networks, and VLANs by assigning different custom network analysis policies to preprocess matching traffic.

Note For an access control policy with rule action as **Trust** and a prefilter rule with action as **Fastpath** with logging options disabled, you will observe that the end-of-flow events are still generated in the system. The events are not visible on the management center event pages.

Access Control Rules: Intrusion Policy Selection

After initial preprocessing, access control rules (when present) evaluate traffic. In most cases, the first access control rule that a packet matches is the rule that handles that traffic; you can monitor, trust, block, or allow matching traffic.

When you allow traffic with an access control rule, the system can inspect the traffic for discovery data, malware, prohibited files, and intrusions, in that order. Traffic not matching any access control rule is handled by the access control policy's default action, which can also inspect for discovery data and intrusions.

Note All packets, **regardless** of which network analysis policy preprocesses them, are matched to configured access control rules—and thus are potentially subject to inspection by intrusion policies—in top-down order.

The diagram in How Policies Examine Traffic For Intrusions, on page 2036 shows the flow of traffic through a device in an inline, intrusion prevention and AMP for Networks deployment, as follows:

- Access Control Rule A allows matching traffic to proceed. The traffic is then inspected for discovery data by the network discovery policy, for prohibited files and malware by File Policy A, and then for intrusions by Intrusion Policy A.
- Access Control Rule B also allows matching traffic. However, in this scenario, the traffic is not inspected for intrusions (or files or malware), so there are no intrusion or file policies associated with the rule. Note that by default, traffic that you allow to proceed is inspected by the network discovery policy; you do not need to configure this.
- In this scenario, the access control policy's default action allows matching traffic. The traffic is then inspected by the network discovery policy, and then by an intrusion policy. You can (but do not have to) use a different intrusion policy when you associate intrusion policies with access control rules or the default action.

The example in the diagram does not include any blocking or trusting rules because the system does not inspect blocked or trusted traffic.

Intrusion Inspection: Intrusion Policies, Rules, and Variable Sets

You can use intrusion prevention as the system's last line of defense before traffic is allowed to proceed to its destination. Intrusion policies govern how the system inspects traffic for security violations and, in inline deployments, can block or alter malicious traffic. The main function of intrusion policies is to manage which intrusion and preprocessor rules are enabled and how they are configured.

Intrusion and Inspector Rules

An intrusion rule is a specified set of keywords and arguments that detects attempts to exploit vulnerabilities on your network; the system uses an intrusion rule to analyze network traffic to check if it matches the criteria in the rule. The system compares packets against the conditions specified in each rule and, if the packet data matches all the conditions specified in a rule, the rule triggers.

The system includes the following types of rules created by Cisco Talos Intelligence Group (Talos):

- *shared object intrusion rules*, which are compiled and cannot be modified (except for rule header information such as source and destination ports and IP addresses)
- *standard text intrusion rules*, which can be saved and modified as new custom instances of the rule.
- *preprocessor rules*, which are rules associated with inspectors and packet decoder detection options in the network analysis policy. You cannot copy or edit inspector rules. Most inspector rules are disabled by default; you must enable them to use inspectors to generate events and, in an inline deployment, drop offending packets.

When the system processes packets according to an intrusion policy, first a rule optimizer classifies all activated rules in subsets based on criteria such as: transport layer, application protocol, direction to or from the protected network, and so on. Then, the intrusion rules engine selects the appropriate rule subsets to apply to each packet. Finally, a multi-rule search engine performs three different types of searches to determine if the traffic matches the rule:

- The protocol field search looks for matches in particular fields in an application protocol.
- The generic content search looks for ASCII or binary byte matches in the packet payload.
- The packet anomaly search looks for packet headers and payloads that, rather than containing specific content, violate well-established protocols.

In a custom intrusion policy, you can tune detection by enabling and disabling rules, as well as by writing and adding your own standard text rules. You can also use Cisco recommendations to associate the operating systems, servers, and client application protocols detected on your network with rules specifically written to protect those assets.

Note When there are insufficient packets to process specific traffic against a block rule, the system continues to evaluate the remaining traffic against other rules. If any of the remaining traffic matches a rule which is set to block, then the session is blocked. However, if the system analyses the remaining traffic to be passed, then the traffic status shows pending on the rule which is stuck for want of complete packets.

Variable Sets

Whenever the system uses an intrusion policy to evaluate traffic, it uses an associated *variable set*. Most variables in a set represent values commonly used in intrusion rules to identify source and destination IP addresses and ports. You can also use variables in intrusion policies to represent IP addresses in rule suppressions and dynamic rule states.

The system provides a single default variable set, which is comprised of predefined default variables. Most system-provided shared object rules and standard text rules use these predefined default variables to define networks and port numbers. For example, the majority of the rules use the variable `$HOME_NET` to specify the protected network and the variable `$EXTERNAL_NET` to specify the unprotected (or outside) network. In addition, specialized rules often use other predefined variables. For example, rules that detect exploits against web servers use the `$HTTP_SERVERS` and `$HTTP_PORTS` variables.

Tip Even if you use system-provided intrusion policies, Cisco **strongly** recommends that you modify key default variables in the default set. When you use variables that accurately reflect your network environment, processing is optimized and the system can monitor relevant systems for suspicious activity. Advanced users can create and use custom variable sets for pairing with one or more custom intrusion policies.

Important If you are creating a custom variable set, do not use a number as the first character in a custom variable set name (for example, 3Snort). This will cause Snort 3 validation to fail when you deploy a configuration to Firewall Threat Defense firewall on the Firewall Management Center.

Intrusion Event Generation

When the system identifies a possible intrusion, it generates an *intrusion or preprocessor event* (sometimes collectively called *intrusion events*). Managed devices transmit their events to the Firewall Management Center, where you can view the aggregated data and gain a greater understanding of the attacks against your network assets. In an inline deployment, managed devices can also drop or replace packets that you know to be harmful.

Each intrusion event in the database includes an event header and contains information about the event name and classification; the source and destination IP addresses; ports; the process that generated the event; and the date and time of the event, as well as contextual information about the source of the attack and its target. For packet-based events, the system also logs a copy of the decoded packet header and payload for the packet or packets that triggered the event.

The packet decoder, the preprocessors, and the intrusion rules engine can all cause the system to generate an event. For example:

- If the packet decoder (configured in the network analysis policy) receives an IP packet that is less than 20 bytes, which is the size of an IP datagram without any options or payload, the decoder interprets this as anomalous traffic. If, later, the accompanying decoder rule in the intrusion policy that examines the packet is enabled, the system generates a inspector event.
- If the IP defragmentation preprocessor encounters a series of overlapping IP fragments, the inspector interprets this as a possible attack and, when the accompanying inspector rule is enabled, the system generates a inspector event.
- Within the intrusion rules engine, most standard text rules and shared object rules are written so that they generate intrusion events when triggered by packets.

As the database accumulates intrusion events, you can begin your analysis of potential attacks. The system provides you with the tools you need to review intrusion events and evaluate whether they are important in the context of your network environment and your security policies.

Asymmetrical Flow Inspection in Snort

In an inline deployment with asymmetrical routing, packet normalization is compromised due to Snort's limited visibility of unidirectional traffic. Snort cannot account for TCP handshake parameters such as window scaling or maximum segment size (MSS) from the unseen flow direction, potentially leading to the host receiving a high volume of packets.

In the following illustration, both devices are running Snort engines. However, neither engine observes the complete traffic flow. The TCP three-way handshake of the flow is not fully captured, which limits the types of normalizations that can be applied. However, other effective normalizations are carried out on the side of the flow that is visible to the Snort engine.

Figure 446: Asymmetric Routing

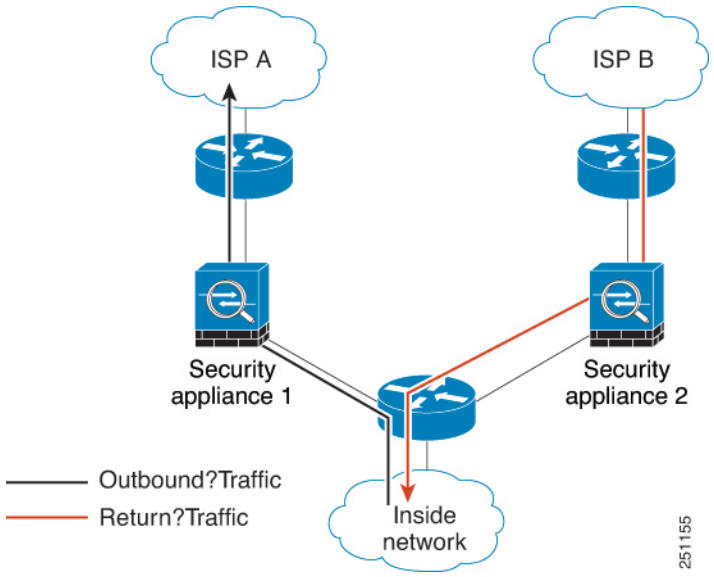

In environments with asymmetric routing, Snort seamlessly adapts to the dynamics without the need for additional configuration. It dynamically adjusts its operations based on flow patterns. Note that asymmetric traffic can potentially impact firewall effectiveness and may not be the optimal choice. However, Snort is designed to provide support for such deployments when necessary.

System-Provided and Custom Network Analysis and Intrusion Policies

Creating a new access control policy is one of the first steps in managing traffic flow using the system. By default, a newly created access control policy invokes system-provided network analysis and intrusion policies to examine traffic.

The following diagram shows how a newly created access control policy in an inline, intrusion-prevention deployment initially handles traffic. The preprocessing and intrusion prevention phases are highlighted.

Note how:

- A default network analysis policy governs the preprocessing of *all* traffic handled by the access control policy. Initially, the system-provided *Balanced Security and Connectivity network analysis policy* is the default.
- The default action of the access control policy allows all non-malicious traffic, as determined by the system-provided *Balanced Security and Connectivity intrusion policy*. Because the default action allows traffic to pass, the discovery feature can examine it for host, application, and user data before the intrusion policy can examine and potentially block malicious traffic.

- The policy uses default Security Intelligence options (global Block and Do Not Block lists only), does not decrypt encrypted traffic with an SSL policy, and does not perform special handling and inspection of network traffic using access control rules.

A simple step you can take to tune your intrusion prevention deployment is to use a different set of system-provided network analysis and intrusion policies as your defaults. Cisco delivers several pairs of these policies with the system.

Or, you can tailor your intrusion prevention deployment by creating and using custom policies. You may find that the inspector options, intrusion rule, and other advanced settings configured in those policies do not address the security needs of your network. By tuning your network analysis and intrusion policies you can configure, at a very granular level, how the system processes and inspects the traffic on your network for intrusions.

System-Provided Network Analysis and Intrusion Policies

Cisco delivers several pairs of network analysis and intrusion policies with the system. By using system-provided network analysis and intrusion policies, you can take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos provides intrusion and inspector rule states as well as initial configurations for inspectors and other advanced settings.

No system-provided policy covers every network profile, traffic mix, or defensive posture. Each covers common cases and network setups that provide a starting point for a well-tuned defensive policy. Although you can use system-provided policies as-is, Cisco strongly recommends that you use them as the base for custom policies that you tune to suit your network.

Tip Even if you use system-provided network analysis and intrusion policies, you should configure the system's intrusion variables to accurately reflect your network environment. At a minimum, modify key default variables in the default set.

As new vulnerabilities become known, Talos releases intrusion rule updates also known as *Lightweight Security Package* (LSP). These rule updates can modify any system-provided network analysis or intrusion policy, and can provide new and updated intrusion rules and inspector rules, modified states for existing rules, and modified default policy settings. Rule updates may also delete rules from system-provided policies and provide new rule categories, as well as modify the default variable set.

If a rule update affects your deployment, the web interface marks affected intrusion and network analysis policies as out of date, as well as their parent access control policies. You must re-deploy an updated policy for its changes to take effect.

For your convenience, you can configure rule updates to automatically re-deploy affected intrusion policies, either alone or in combination with affected access control policies. This allows you to easily and automatically keep your deployment up-to-date to protect against recently discovered exploits and intrusions.

To ensure up-to-date preprocessing settings, you **must** re-deploy access control policies, which also deploys any associated SSL, network analysis, and file policies that are different from those currently running, and can also can update default values for advanced preprocessing and performance options.

Cisco delivers the following network analysis and intrusion policies with the system:

Balanced Security and Connectivity network analysis and intrusion policies

These policies are built for both speed and detection. Used together, they serve as a good starting point for most organizations and deployment types. The system uses the Balanced Security and Connectivity policies and settings as defaults in most cases.

Connectivity Over Security network analysis and intrusion policies

These policies are built for organizations where connectivity (being able to get to all resources) takes precedence over network infrastructure security. The intrusion policy enables far fewer rules than those enabled in the Security over Connectivity policy. Only the most critical rules that block traffic are enabled.

Security Over Connectivity network analysis and intrusion policies

These policies are built for organizations where network infrastructure security takes precedence over user convenience. The intrusion policy enables numerous network anomaly intrusion rules that could alert on or drop legitimate traffic.

Maximum Detection network analysis and intrusion policies

These policies are built for organizations where network infrastructure security is given even more emphasis than is given by the Security Over Connectivity policies, with the potential for even greater operational impact. For example, the intrusion policy enables rules in a large number of threat categories including malware, exploit kit, old and common vulnerabilities, and known in-the-wild exploits.

No Rules Active intrusion policy

In the No Rules Active intrusion policy, all intrusion rules, and all advanced settings except intrusion rule thresholds, are disabled. This policy provides a starting point if you want to create your own intrusion policy instead of basing it on the enabled rules in one of the other system-provided policies.

Note Depending on the system-provided base policy that is selected, the settings of the policy vary. To view the policy settings, click the **Edit** icon next to the policy and then click the **Base Policy** drop-down box.

Benefits of Custom Network Analysis and Intrusion Policies

You may find that the inspector options, intrusion rules, and other advanced settings configured in the system-provided network analysis and intrusion policies do not fully address the security needs of your organization.

Building custom policies can improve the performance of the system in your environment and can provide a focused view of the malicious traffic and policy violations occurring on your network. By creating and tuning custom policies you can configure, at a very granular level, how the system processes and inspects the traffic on your network for intrusions.

All custom policies have a base policy, also called a base layer, which defines the default settings for all configurations in the policy. A layer is a building block that you can use to efficiently manage multiple network analysis or intrusion policies.

In most cases, you base custom policies on system-provided policies, but you can use another custom policy. However, all custom policies have a system-provided policy as the eventual base in a policy chain. Because rule updates can modify system-provided policies, importing a rule update may affect you even if you are using a custom policy as your base. If a rule update affects your deployment, the web interface marks affected policies as out of date.

Benefits of Custom Network Analysis Policies

By default, one network analysis policy preprocesses all unencrypted traffic handled by the access control policy. That means that all packets are decoded and preprocessed according to the same settings, regardless of the intrusion policy (and therefore intrusion rule set) that later examines them.

Initially, the system-provided Balanced Security and Connectivity network analysis policy is the default. A simple way to tune preprocessing is to create and use a custom network analysis policy as the default.

Tuning options available vary by inspector, but some of the ways you can tune inspectors and decoders include:

- You can disable inspectors that do not apply to the traffic you are monitoring. For example, the HTTP Inspect inspector normalizes HTTP traffic. If you are confident that your network does not include any web servers using Microsoft Internet Information Services (IIS), you can disable the inspector option that looks for IIS-specific traffic and thereby reduce system processing overhead.

Note If you disable a inspector in a custom network analysis policy, but the system needs to use that inspector to later evaluate packets against an enabled intrusion or inspector rule, the system automatically enables and uses the inspector although the inspector remains disabled in the network analysis policy web interface.

- Specify ports, where appropriate, to focus the activity of certain inspectors. For example, you can identify additional ports to monitor for DNS server responses or encrypted SSL sessions, or ports on which you decode telnet, HTTP, and RPC traffic.

For advanced users with complex deployments, you can create multiple network analysis policies, each tailored to preprocess traffic differently. Then, you can configure the system to use those policies to govern the preprocessing of traffic using different security zones, networks, or VLANs. (Note that ASA FirePOWER modules cannot restrict preprocessing by VLAN.)

Note Tailoring preprocessing using custom network analysis policies—especially multiple network analysis policies—is an advanced task. Because preprocessing and intrusion inspection are so closely related, you **must** be careful to allow the network analysis and intrusion policies examining a single packet to complement each other.

Benefits of Custom Intrusion Policies

In a newly created access control policy initially configured to perform intrusion prevention, the default action allows all traffic, but first inspects it with the system-provided Balanced Security and Connectivity intrusion policy. Unless you add access control rules or change the default action, all traffic is inspected by that intrusion policy.

To customize your intrusion prevention deployment, you can create multiple intrusion policies, each tailored to inspect traffic differently. Then, configure an access control policy with rules that specify which policy inspects which traffic. Access control rules can be simple or complex, matching and inspecting traffic using multiple criteria including security zone, network or geographical location, VLAN, port, application, requested URL, or user.

The main function of intrusion policies is to manage which intrusion and inspector rules are enabled and how they are configured, as follows:

- Within each intrusion policy, you should verify that all rules applicable to your environment are enabled, and improve performance by disabling rules that are not applicable to your environment. You can specify which rules should drop or modify malicious packets.
- Cisco recommendations allow you to associate the operating systems, servers, and client application protocols detected on your network with rules specifically written to protect those assets.
- You can modify existing rules and write new standard text rules as needed to catch new exploits or to enforce your security policies.

Other customizations you might make to an intrusion policy include:

- The sensitive data preprocessor detects sensitive data such as credit card numbers and Social Security numbers in ASCII text. Note that other inspectors that detect specific threats (back orifice attacks, several portscan types, and rate-based attacks that attempt to overwhelm your network with excessive traffic) are configured in network analysis policies.
- Global thresholds cause the system to generate events based on how many times traffic matching an intrusion rule originates from or is targeted to a specific address or address range within a specified time period. This helps prevent the system from being overwhelmed with a large number of events.
- Suppressing intrusion event notifications and setting thresholds for individual rules or entire intrusion policies can also can prevent the system from being overwhelmed with a large number of events.
- In addition to the various views of intrusion events within the web interface, you can enable logging to syslog facilities or send event data to an SNMP trap server. Per policy, you can specify intrusion event notification limits, set up intrusion event notification to external logging facilities, and configure external responses to intrusion events. Note that in addition to these per-policy alerting configurations, you can globally enable or disable email alerting on intrusion events for each rule or rule group. Your email alert settings are used regardless of which intrusion policy processes a packet.

Limitations of Custom Policies

Because preprocessing and intrusion inspection are so closely related, you **must** be careful that your configuration allows the network analysis and intrusion policies processing and examining a single packet to complement each other.

By default, the system uses one network analysis policy to preprocess all traffic handled by managed devices using a single access control policy. The following diagram shows how a newly created access control policy in an inline, intrusion-prevention deployment initially handles traffic. The preprocessing and intrusion prevention phases are highlighted.

Notice how a default network analysis policy governs the preprocessing of *all* traffic handled by the access control policy. Initially, the system-provided Balanced Security and Connectivity network analysis policy is the default.

A simple way to tune preprocessing is to create and use a custom network analysis policy as the default. However, if you disable a inspector in a custom network analysis policy but the system needs to evaluate preprocessed packets against an enabled intrusion or inspector rule, the system automatically enables and uses the inspector although it remains disabled in the network analysis policy web interface.

Note In order to get the performance benefits of disabling an inspector, you **must** make sure that none of your intrusion policies have enabled rules that require that inspector.

An additional challenge arises if you use multiple custom network analysis policies. For advanced users with complex deployments, you can tailor preprocessing to specific security zones, networks, and VLANs by assigning custom network analysis policies to preprocess matching traffic. (Note that ASA FirePOWER cannot restrict preprocessing by VLAN.) To accomplish this, you add custom *network analysis rules* to your access control policy. Each rule has an associated network analysis policy that governs the preprocessing of traffic that matches the rule.

Tip You configure network analysis rules as an advanced setting in an access control policy. Unlike other types of rules, network analysis rules invoke—rather than being contained by—network analysis policies.

The system matches packets to any configured network analysis rules in top-down order by rule number. Traffic that does not match any network analysis rule is preprocessed by the default network analysis policy. While this allows you a great deal of flexibility in preprocessing traffic, keep in mind that all packets, **regardless** of which network analysis policy preprocessed them, are subsequently matched to access control rules—and thus to potential inspection by intrusion policies—in their own process. In other words, preprocessing a packet with a particular network analysis policy does **not** guarantee that the packet will be examined with any particular intrusion policy. You **must** carefully configure your access control policy so it invokes the correct network analysis and intrusion policies to evaluate a particular packet.

The following diagram shows in focused detail how the network analysis policy (preprocessing) selection phase occurs before and separately from the intrusion prevention (rules) phase. For simplicity, the diagram eliminates the discovery and file/malware inspection phases. It also highlights the default network analysis and default-action intrusion policies.

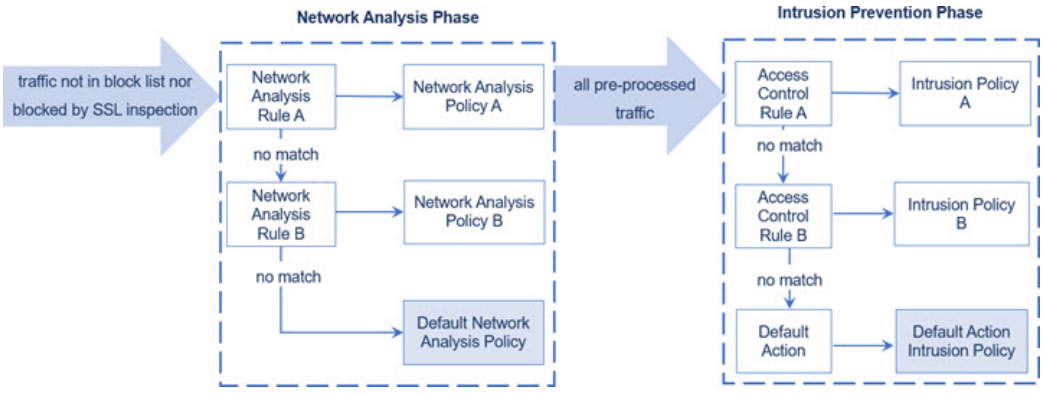

In this scenario, an access control policy is configured with two network analysis rules and a default network analysis policy:

- Network Analysis Rule A preprocesses matching traffic with Network Analysis Policy A. Later, you want this traffic to be inspected by Intrusion Policy A.
- Network Analysis Rule B preprocesses matching traffic with Network Analysis Policy B. Later, you want this traffic to be inspected by Intrusion Policy B.

- All remaining traffic is preprocessed with the default network analysis policy. Later, you want this traffic to be inspected by the intrusion policy associated with the access control policy's default action.

After the system preprocesses traffic, it can examine the traffic for intrusions. The diagram shows an access control policy with two access control rules and a default action:

- Access Control Rule A allows matching traffic. The traffic is then inspected by Intrusion Policy A.
- Access Control Rule B allows matching traffic. The traffic is then inspected by Intrusion Policy B.
- The access control policy's default action allows matching traffic. The traffic is then inspected by the default action's intrusion policy.

Each packet's handling is governed by a network analysis policy and intrusion policy pair, but the system does **not** coordinate the pair for you. Consider a scenario where you misconfigure your access control policy so that Network Analysis Rule A and Access Control Rule A do not process the same traffic. For example, you could intend the paired policies to govern the handling of traffic on a particular security zone, but you mistakenly use different zones in the two rules' conditions. This could cause traffic to be incorrectly preprocessed. For this reason, tailoring preprocessing using network analysis rules and custom policies is an **advanced** task.

Note that for a single connection, although the system selects a network analysis policy before an access control rule, some preprocessing (notably application layer preprocessing) occurs after access control rule selection. This does **not** affect how you configure preprocessing in custom network analysis policies.

Prerequisites for Network Analysis and Intrusion Policies

To allow the Snort inspection engine to process traffic for intrusion and malware analysis, you must have the IPS license enabled for the Firewall Threat Defense device.

You must be an Admin user to manage network analysis, intrusion policies, and perform migration tasks.

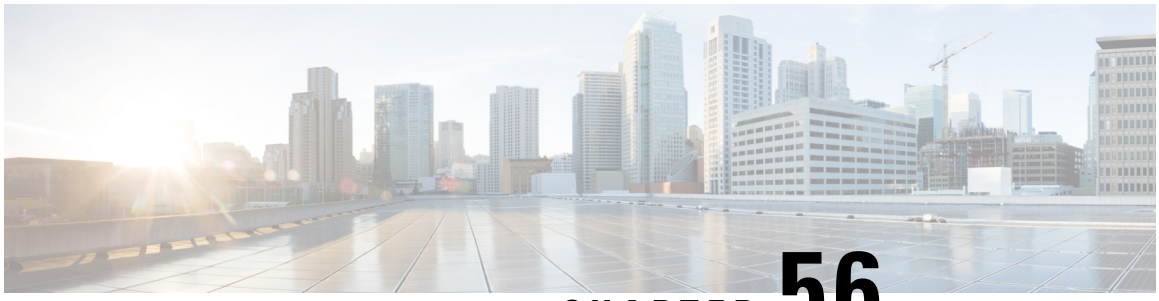

CHAPTER 56

Migrate from Snort 2 to Snort 3

Starting with Version 7.0, Snort 3 is the default inspection engine for new Firewall Threat Defense deployments with Firewall Management Center. Snort 3 is required in Version 7.7+. If you are still using the Snort 2 inspection engine, switch to Snort 3 now for improved detection and performance.

Upgrading Firewall Threat Defense to Version 7.2 through 7.6 also upgrades eligible Snort 2 devices to Snort 3. For devices that are ineligible because they use custom intrusion or network analysis policies, manually upgrade to Snort 3 as described here.

Although you can switch individual devices back, you should not. Upgrading Firewall Threat Defense to Version 7.7+ is blocked for Snort 2 devices.

- Snort 3 Inspection Engine, on page 2049
- Prerequisites for Network Analysis and Intrusion Policies, on page 2050
- How to Migrate from Snort 2 to Snort 3, on page 2050
- View Snort 2 and Snort 3 Base Policy Mapping, on page 2054
- Synchronize Snort 2 Rules with Snort 3 , on page 2054
- Deploy Configuration Changes, on page 2055

Snort 3 Inspection Engine

Snort 3 is the default inspection engine for newly registered Firewall Threat Defense devices of version 7.0 and later. However, for Firewall Threat Defense devices of lower versions, Snort 2 is the default inspection engine. When you upgrade a managed Firewall Threat Defense device to version 7.0 or later, the inspection engine remains on Snort 2. To use Snort 3 in upgraded Firewall Threat Defenses of version 7.0 and later, you must explicitly enable it. When Snort 3 is enabled as the inspection engine of the device, the Snort 3 version of the intrusion policy that is applied on the device (through the access control policies) is activated and applied to all the traffic passing through the device.

You can switch Snort versions when required. Snort 2 and Snort 3 intrusion rules are mapped and the mapping is system-provided. However, you may not find a one-to-one mapping of all the intrusion rules in Snort 2 and Snort 3. If you change the rule action for one rule in Snort 2, that change is not retained if you switch to Snort 3 without synchronizing Snort 2 with Snort 3. For more information on synchronization, see Synchronize Snort 2 Rules with Snort 3 , on page 2054.

Prerequisites for Network Analysis and Intrusion Policies

To allow the Snort inspection engine to process traffic for intrusion and malware analysis, you must have the IPS license enabled for the Firewall Threat Defense device.

You must be an Admin user to manage network analysis, intrusion policies, and perform migration tasks.

How to Migrate from Snort 2 to Snort 3

Migrating from Snort 2 to Snort 3 requires you to switch the inspection engine of the Firewall Threat Defense device from Snort 2 to Snort 3.

Depending on your requirements, the tasks to complete the migration of your device from Snort 2 to Snort 3 is listed in the following table:

Step	Task	Links to Procedures
1	Enable Snort 3	• Enable Snort 3 on an Individual Device, on page 2051 • Enable Snort 3 on Multiple Devices, on page 2051
2	Convert Snort 2 custom rules to Snort 3	• Convert all Snort 2 Custom Rules across all Intrusion Policies to Snort 3, on page 2052 • Convert Snort 2 Custom Rules of a Single Intrusion Policy to Snort 3, on page 2053
3	Synchronize Snort 2 rules with Snort 3	Synchronize Snort 2 Rules with Snort 3 , on page 2054

Prerequisites for Migrating from Snort 2 to Snort 3

The following are the recommended prerequisites that you must consider before migrating your device from Snort 2 to Snort 3.

- Have a working knowledge of Snort. To learn about the Snort 3 architecture, see Snort 3 Adoption.
- Back up your management center. See Backup the Management Center.
- Back up your intrusion policy. See Exporting Configurations.
- Clone your intrusion policy. To do this, you can use an existing policy as the base policy to create a copy of your intrusion policy. In the **Intrusion Policies** page, click **Create Policy** and choose an existing intrusion policy from the **Base Policy** dropdown list.

Enable Snort 3 on an Individual Device

Important During the deployment process, there could be a momentary traffic loss because the current inspection engine needs to be shut down.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click the device to go to the device home page.

Note
The device is marked as Snort 2 or Snort 3, showing the current version on the device.

Step 3 Click the **Device** tab.

Step 4 In the Inspection Engine section, click **Upgrade**.

Note
In case you want to disable Snort 3, click **Revert to Snort 2** in the Inspection Engine section.

Step 5 Click **Yes**.

What to do next

Deploy the changes on the device. See, Deploy Configuration Changes, on page 2055.

The system converts your policy configurations during the deployment process to make them compatible with the selected Snort version.

Enable Snort 3 on Multiple Devices

To enable Snort 3 on multiple devices, ensure all the required Firewall Threat Defense devices are on version 7.0 or later.

Important During the deployment process, there could be a momentary traffic loss because the current inspection engine needs to be shut down.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Select all the devices on which you want to enable or disable Snort 3.

Note

The devices are marked as Snort 2 or Snort 3, showing the current version on the device.

Step 3 Click **Select Bulk Action** drop-down list and choose **Upgrade to Snort 3**.

Step 4 Click **Yes**.

What to do next

Deploy the changes on the device. See, Deploy Configuration Changes, on page 2055.

The system converts your policy configurations during the deployment process to make them compatible with the selected Snort version.

Convert Snort 2 Custom IPS Rules to Snort 3

If you are using a rule set from a third-party vendor, contact that vendor to confirm that their rules successfully convert to Snort 3 or to obtain a replacement rule set written natively for Snort 3. If you have custom rules that you have written yourself, familiarize with writing Snort 3 rules prior to conversion, so you can update your rules to optimize Snort 3 detection after conversion. See the links below to learn more about writing rules in Snort 3.

- https://blog.snort.org/2020/08/how-rules-are-improving-in-snort-3.html
- https://blog.snort.org/2020/10/talos-transition-to-snort-3.html

You can refer to other blogs at https://blog.snort.org/ to learn more about Snort 3 rules.

See the following procedures to convert Snort 2 rules to Snort 3 rules using the system-provided tool.

- Convert all Snort 2 Custom Rules across all Intrusion Policies to Snort 3, on page 2052
- Convert Snort 2 Custom Rules of a Single Intrusion Policy to Snort 3, on page 2053

Important Snort 2 network analysis policy (NAP) settings *cannot* be copied to Snort 3 automatically. NAP settings have to be manually replicated in Snort 3.

Convert all Snort 2 Custom Rules across all Intrusion Policies to Snort 3

Procedure

Step 1 Choose **Objects** > **Intrusion Rules** .

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Ensure **All Rules** is selected in the left pane.

Step 4 Click the **Tasks** drop-down list and choose:

- **Convert Snort 2 rules and import**—To automatically convert all the Snort 2 custom rules across all the intrusion policies to Snort 3 and import them into Firewall Management Center as Snort 3 custom rules.

- **Convert Snort 2 rules and download**—To automatically convert all the Snort 2 custom rules across all the intrusion policies to Snort 3 and download them into your local system.

Step 5 Click **OK**.

Note

- If you selected **Convert and import** in the previous step, then all the converted rules are saved under a newly created rule group **All Snort 2 Converted Global** under **Local Rules**.
- If you selected **Convert and download** in the previous step, then save the rules file locally. You can review the converted rules in the downloaded file and later upload them by following the steps in Add Custom Rules to Rule Groups, on page 2087.

Refer to the video Converting Snort 2 Rules to Snort 3 for additional support and information.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

Convert Snort 2 Custom Rules of a Single Intrusion Policy to Snort 3

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 In the **Intrusion Policies** tab, click **Show Snort 3 Sync status**.

Step 3 Click the **Sync** icon **Snort out-of-Sync** () of the intrusion policy.

Note
If the Snort 2 and the Snort 3 versions of the intrusion policy are synchronized, then the **Sync** icon is in blue**Snort in-Sync** (). It indicates that there are no custom rules to be converted.

Step 4 Read through the summary and click the **Custom Rules** tab.

Step 5 Choose:

- **Import converted rules to this policy**—To convert the Snort 2 custom rules in the intrusion policy to Snort 3 and import them into Firewall Management Center as Snort 3 custom rules.
- **Download converted rules**—To convert the Snort 2 custom rules in the intrusion policy to Snort 3 and download them into your local system. You can review the converted rules in the downloaded file and later upload the file by clicking the upload icon.

Step 6 Click **Re-Sync**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

View Snort 2 and Snort 3 Base Policy Mapping

Note Snort 2 is not supported on threat defense Version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the Firewall Management Center guide that matches your Firewall Threat Defense version.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Ensure the **Intrusion Policies** tab is selected.

Step 3 Click **IPS Mapping**.

Step 4 In the **IPS Policy Mapping** dialog box, click **View Mappings** to view the Snort 3 to Snort 2 intrusion policy mapping.

Step 5 Click **OK**.

Synchronize Snort 2 Rules with Snort 3

To ensure that the Snort 2 version settings and custom rules are retained and carried over to Snort 3, the Firewall Management Center provides the synchronization functionality. Synchronization helps Snort 2 rule override settings and custom rules, which you may have altered and added over the last few months or years, to be replicated on the Snort 3 version. This utility helps to synchronize Snort 2 version policy configuration with Snort 3 version to start with similar coverage.

Note Snort 2 is not supported on threat defense Version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the Firewall Management Center guide that matches your Firewall Threat Defense version.

If the Firewall Management Center is upgraded from 6.7 or earlier to 7.0 or later version, the system synchronizes the configuration. If the Firewall Management Center is a fresh 7.0 or later version, you can upgrade to a higher version, and the system will not synchronize any content during upgrade.

Before upgrading a device to Snort 3, if changes are made in Snort 2 version, you can use this utility to have the latest synchronization from Snort 2 version to Snort 3 version so that you start with a similar coverage.

Note On moving to Snort 3, it is recommended that you manage the Snort 3 version of the policy independently and do not use this utility as a regular operation.

Important

- Only the Snort 2 rule overrides and custom rules are copied to Snort 3 and not the other way around. You may not find a one-to-one mapping of all the intrusion rules in Snort 2 and Snort 3. Your changes to rule actions for rules that exist in both versions are synchronized when you perform the following procedure.
- Synchronization *does not* migrate the threshold and suppression settings of any custom or system-provided rules from Snort 2 to Snort 3.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Ensure the **Intrusion Policies** tab is selected.

Step 3 Click **Show Snort 3 Sync status**.

Step 4 Identify the intrusion policy that is out-of-sync.

Step 5 Click the **Sync** icon **Snort out-of-Sync** ( ).

Note
If the Snort 2 and the Snort 3 versions of the intrusion policy are synchronized, then the **Sync** icon is in blue**Snort in-Sync** ( ).

Step 6 Read through the summary and download a copy of the summary if required.

Step 7 Click **Re-Sync**.

Note
- The synchronized settings will be applicable on the Snort 3 intrusion engine only if it is applied on a device, and after a successful deployment.
- Snort 2 custom rules can be converted to Snort 3 using the system-provided tool. If you have any Snort 2 custom rules click the Custom Rules tab and follow the on-screen instructions to convert the rules. For more information, see Convert Snort 2 Custom Rules of a Single Intrusion Policy to Snort 3, on page 2053.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

Deploy Configuration Changes

After you change configurations, deploy them to the affected devices.

Note This topic covers the basic steps involved in deploying configuration changes. We *strongly* recommend that you refer the *Deploy Configuration Changes* topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide* to understand the prerequisites and implications of deploying the changes before proceeding with the steps.

Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic.

Procedure

Step 1 On the Secure Firewall Management Center menu bar, click **Deploy** and choose **Deployment**.

The GUI page lists the devices with out-of-date configurations having **Pending**status.

- The **Modified By** column lists the users who have modified the policies or objects. Expand the device listing to view the users who have modified the policies for each policy listing.

  Note
  Usernames are not provided for deleted policies and objects.

- The **Inspect Interruption** column indicates if traffic inspection interruption might occur in the device during deployment.

  If this column is blank for a device, it indicates that there will be no traffic inspection interruptions on that device during deployment.

- The **Last Modified Time** column specifies the last time you made configuration changes.
- The **Preview** column allows you to preview the changes for the next deployment.
- The **Status** column provides the status for each deployment.

Step 2 Identify and choose the devices on which you want to deploy configuration changes.

- Search—Search for the device name, type, domain, group, or status in the search box.
- Expand—Click **Expand Arrow** (>) to view device-specific configuration changes to be deployed.

  When you check a check box adjacent to a device, all the changes made to the device and listed under the device, are pushed for deployment. However, you can use **Policy selection** ( ) to select individual policies or specific configurations to deploy while withholding the remaining changes without deploying them.

  Note
  - When the status in the **Inspect Interruption** column indicates (**Yes**) that deploying will interrupt inspection, and perhaps traffic, on a Firewall Threat Defense device, the expanded list indicates the specific configurations causing the interruption with the **Inspect Interruption** ( ).

- When there are changes to interface groups, security zones, or objects, the impacted devices are shown as out-of-date on the Firewall Management Center. To ensure that these changes take effect, the policies with these interface groups, security zones, or objects, also need to be deployed along with these changes. The impacted policies are shown as out-of-date on the **Preview**page on the Firewall Management Center.

Step 3 Click **Deploy**.

Step 4 If the system identifies errors or warnings in the changes to be deployed, it displays them in the **Validation Messages** window. To view complete details, click the arrow icon before the warnings or errors.

You have the following choices:

- Deploy—Continue deploying without resolving warning conditions. You cannot proceed if the system identifies errors.
- Close—Exit without deploying. Resolve the error and warning conditions, and attempt to deploy the configuration again.

What to do next

During deployment, if there is a deployment failure, there is a possibility that the failure may impact traffic. However, it depends on certain conditions. If there are specific configuration changes in the deployment, the deployment failure may lead to traffic being interrupted. For details, see the Deploy Configuration Changes topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide*.

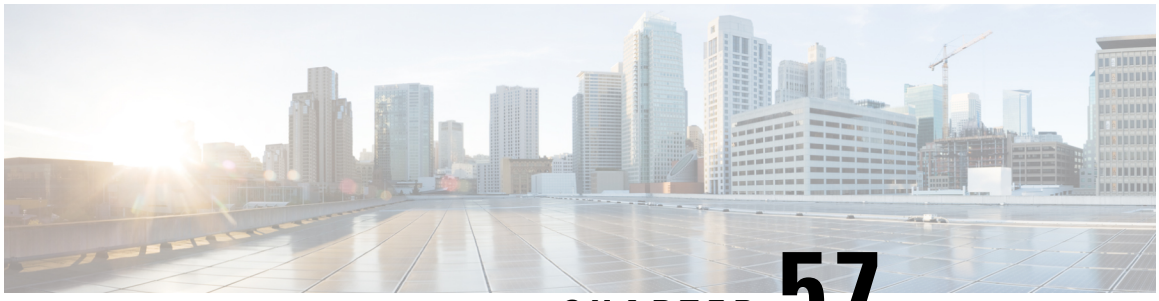

CHAPTER 57

Get Started with Snort 3 Intrusion Policies

This chapter provides information on managing Snort 3 intrusion policies and access control rule configurations for intrusion detection and prevention.

- Overview of Intrusion Policies, on page 2059
- Prerequisites for Network Analysis and Intrusion Policies, on page 2060
- Create a Custom Snort 3 Intrusion Policy , on page 2060
- Edit Snort 3 Intrusion Policies, on page 2061
- Change the Base Policy of an Intrusion Policy, on page 2066
- View Snort 2 and Snort 3 Base Policy Mapping, on page 2066
- Synchronize Snort 2 Rules with Snort 3 , on page 2067
- Manage Intrusion Policies, on page 2068
- Access Control Rule Configuration to Perform Intrusion Prevention, on page 2069
- Deploy Configuration Changes, on page 2071

Overview of Intrusion Policies

Intrusion policies are defined sets of intrusion detection and prevention configurations that inspect traffic for security violations and, in inline deployments, can block or alter malicious traffic. Intrusion policies are invoked by your access control policy and are the system's last line of defense before traffic is allowed to its destination.

At the heart of each intrusion policy are the intrusion rules. An enabled rule causes the system to generate intrusion events for (and optionally block) traffic matching the rule. Disabling a rule stops processing of the rule.

The system delivers several base intrusion policies, which enable you to take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos sets intrusion and inspector rule states (enabled or disabled), as well as provides the initial configurations for other advanced settings.

Tip System-provided intrusion and network analysis policies are similarly named but contain different configurations. For example, the Balanced Security and Connectivity network analysis policy and the Balanced Security and Connectivity intrusion policy work together and can both be updated in intrusion rule updates. However, the network analysis policy governs mostly preprocessing options, whereas the intrusion policy governs mostly intrusion rules.

If you create a custom intrusion policy, you can:

- Tune detection by enabling and disabling rules, as well as by writing and adding your own rules.
- Use Secure Firewall recommendations to associate the operating systems, servers, and client application protocols detected on your network with rules specifically written to protect those assets.

An intrusion policy can drop matching packets and generate intrusion events. To configure an intrusion or preprocessor drop rule, set its state to Block.

When tailoring your intrusion policy, especially when enabling and adding rules, keep in mind that some intrusion rules require that traffic first be decoded or preprocessed in a certain way. Before an intrusion policy examines a packet, the packet is preprocessed according to configurations in a network analysis policy. If you disable a required inspector, the system automatically uses it with its current settings, although the inspector remains disabled in the network analysis policy web interface.

Caution Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet **must** complement each other. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task.

After you configure a custom intrusion policy, you can use it as part of your access control configuration by associating the intrusion policy with one or more access control rules or an access control policy's default action. This forces the system to use the intrusion policy to examine certain allowed traffic before the traffic passes to its final destination. A variable set that you pair with the intrusion policy allows you to accurately reflect your home and external networks and, as appropriate, the servers on your network.

Note that by default, the system disables intrusion inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion inspection configured.

Refer to the video for additional support and information - Snort 3 Intrusion Policy Overview.

Prerequisites for Network Analysis and Intrusion Policies

To allow the Snort inspection engine to process traffic for intrusion and malware analysis, you must have the IPS license enabled for the Firewall Threat Defense device.

You must be an Admin user to manage network analysis, intrusion policies, and perform migration tasks.

Create a Custom Snort 3 Intrusion Policy

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Click **Create Policy**.

Step 3 Enter a unique **Name** and, optionally, a **Description**.

Step 4 Choose the **Inspection Mode**.

The selected action determines whether intrusion rules block and alert (**Prevention** mode) or only alert (**Detection** mode).

Note
Before selecting the prevention mode, you might want block rules to alert only so you can identify rules that cause a lot of false positives.

Step 5 Choose the **Base Policy**.

You can use either a system-provided policy or an existing policy as your base policy.

Step 6 Click **Save**.

The new policy has the same settings as its base policy.

What to do next

To customize the policy, see Edit Snort 3 Intrusion Policies, on page 2061.

Edit Snort 3 Intrusion Policies

While editing a Snort 3 policy, all the changes are saved instantaneously. No additional action is required to save the changes.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Ensure the **Intrusion Policies** tab is selected.

Step 3 Click **Snort 3 Version** next to the intrusion policy you want to configure.

Step 4 Edit your policy:

- Change the mode—Click the **Mode** drop-down to change the inspection mode.

Caution
The inspection mode is changed only for the Snort 3 version of the policy. The existing inspection mode is retained in the Snort 2 version as is, which means that your Snort 2 and Snort 3 versions of the policy will have different inspection modes. We recommend you to use this option with caution.

- **Prevention**—Triggered Block rules create an event (alert) and drop the connection.
- **Detection**—Triggered Block rules create an alert.

  You can choose the detection mode before going for prevention. For example, before choosing the prevention mode, you might want block rules to alert only, so that you can identify rules that cause a lot of false positives.

Step 5 Click the **Base Policy** layer that defines the intrusion policy's default settings.

- Search rules—Use the search field to filter the display. You can enter the GID, SID, rule message, or reference info. For example, GID:1; SID:9621—to display only rule 1:962, SID:9621,9622,9623—to display multiple rules with different SIDs. You can also click inside the Search text box to choose any of the following options:
  - apply the filters **Action = Alert**, or **Action: Block**
  - apply the **Disabled Rules** filter
  - show **Custom/User Defined Rules**
  - filter by GID, SID, or GID:SID
  - filter by CVE
  - filter by comment
- View filtered rules—Click any of the **Presets** to view rules that are set to alert, block, disabled, and so on.

  Overridden rules indicate the rules where the rule action has been changed from the default action to a different action. Note that, once changed, the rule action status is Overridden even if you change it back to its original default action. However, if you select **Revert to default** from the **Rule Action** drop-down list, the Overridden status is removed.

  Advanced Filters provides filter options based on the Lightweight Security Package (LSP) releases, Classifications of intrusions, and Microsoft Vulnerabilities.
- View rule documentation—Click the rule ID or the **Rule Documentation** icon to display Talos documentation for the rule.
- View a rule details—Click the **Expand Arrow** (›) icon in a rule row to view the rule details.
- Add rule comments—Click **Comment** (💬) under the Comments column to add comments for a rule.

Step 6 **Group Overrides**—Click the **Group Overrides** layer that lists all the categories of rule groups. The top level parent rule groups with Description, Overrides and Enabled Groups, and so on is displayed. Parent rule groups cannot be updated and are read-only. Only the leaf rule groups can be updated. In each rule group, you can traverse up to the last leaf group. Across each group, you can override, include, and exclude rule groups. In the leaf rule groups, you can:

- Search rule groups—Use the search field to enter keywords and search for rule groups.
- In the left panel, you can choose any of the preset filter options to search for rule groups:
  - All—For displaying all rule groups.
  - Excluded—For excluded groups.
  - Included—For included groups.
  - Overridden—For rule group configuration that is overridden.
- Set the security level for a rule group—Navigate to the required rule group on the left pane and click it. Click **Edit** next to the **Security Level** of the rule group to increase or decrease the security level based on system-defined rule settings.

  In the **Edit Security Level** dialog box, you have the option to click **Revert to Default**, which reverts the changes you made.

The Firewall Management Center automatically changes the action for the rules of the rule group for the configured security level. In the **Rule Overrides** layer, notice the count of Block Rules and Disabled Rules in the **Presets** every time you change the security level.

- You can make bulk changes to the security level to change the security level of all rule groups within a particular rule category. Bulk security level applies to rule groups that have more than one rule group. After a bulk update of rule groups, you can still update the security level of any of the associated rule groups within it.

  There can be **mixed** security levels within rule groups; **mixed** indicates that the child groups contain a mix of security levels within the parent rule group.

- Include or exclude rule groups—The rule groups displayed are the default rule groups associated with the system-provided base intrusion policy. You can include and exclude rule groups from the intrusion policy. An excluded rule group is removed from the intrusion policy and its rules are not applied on the traffic. For information on uploading custom rules in Firewall Management Center, see Add Custom Rules to Rule Groups, on page 2087.

  To exclude a rule group:

  a. Navigate the Rule Groups pane and choose the rule group that you want to exclude.

  b. Click the **Exclude** hyperlink on the right-pane.

  c. Click **Exclude**.

  To include a new rule group or multiple rule groups with the uploaded custom rules or a previously excluded rule group:

  a. Click **Add** (+) next to the rule group filter dropdown list.

  b. Choose all the rule groups you want to add by checking the check box next to it.

  c. Click **Save**.

- For a leaf rule group, click the icon under the **Override** column header to see the rule action trail, which describes the sequence of overridden rule actions that can be assigned due to the base policy and group overrides for an intrusion rule. Rule actions can be obtained from either the base policy configurations or the user group override. The user group override takes the priority between the two; priority refers to the final overridden action that is assigned to the rule group.

- Click the rule count (number) under the **Rule Count** column header to see a summary of rules that are part of the rule group.

Step 7 **Recommendations**—Click the **Recommendations** layer if you want to generate and apply Cisco recommended rules. Recommendations use the host database to enable or disable rules, based on known vulnerabilities.

Step 8 **Rule Overrides**—Click the **Rule Overrides** layer to choose any of the presets to view rules, which are set to alert, block, disabled, overridden, rewrite, pass, drop, or reject.

- The **Set By** column shows the default set by state (Base Policy) or modified rule state by Group Overrides, Rule Overrides, or Recommendations. The **Set By** column in **All Rules** (in the left pane) shows the trail of rule action override actions based on priority order. The priority order of rule actions is Rule Override > Recommendations > Group Override > Base Policy.

- Modify **Rule Action**—To modify rule actions, choose either of the following:

- Bulk edit—Choose one or more rules, then choose the required action from the **Rule Action** drop-down list; and click **Save**.

  Note
  Bulk rule action changes are supported only for the first 500 rules.

- Single rule edit—Choose the action for the rule from the drop-down list in the **Rule Action** column.

Rule actions are:

- **Block**— Generates event, blocks current matching packet and all the subsequent packets in this connection.
- **Alert**— Generates only events for matching packet and does not drop packet or connection.
- **Disable**—Does not match traffic against this rule. No events are generated.
- **Revert to default**—Reverts to the system default action.
- **Pass**— No events are generated, allows packet to pass without further evaluation by any subsequent Snort rules.

  Note
  The Pass action is available only for custom rules and not for system-provided rules.

- **Drop**— Generates event, drops matching packet and does not block further traffic in this connection.
- **Reject**— Generates event, drops matching packet, blocks further traffic in this connection and sends TCP reset if it is a TCP protocol to source and destination hosts.

  Behavior of reject in different firewall modes and IP address or source or destination in relation to Client or Server: Snort sends RST packets to both client and server in cases of routed, inline, and bridged interfaces. Snort sends two RST packets. RST packet in clients directions will have source set to server's IP and destination set to client's IP. RST packet in servers direction will have source set to client's IP and destination set to server's IP.

- **Rewrite**— Generates event and overwrites packet contents based on the replace option in the rule.

For IPS rule action logging, see Rule Action Logging, on page 2065.

If there is a **React** rule, it is converted to an alert action.

Step 9 Click the **Summary** layer for a holistic view of the current changes to the policy. The policy summary page contains the following information:

- Rule distribution of the policy, that is, active rules, disabled rules, and so on.
- Option to export policy and generate report of the intrusion policy.
- Base policy details.
- Option to generate recommendations.
- Group overrides that shows the list of groups that you have overridden.
- Rule overrides that shows the list of rules that you have overridden.

- In the **Summary** layer, click the ? icon to open a popup window of the Snort helper guide that explains the Snort layering concepts.

To change the base policy, see Change the Base Policy of an Intrusion Policy, on page 2066.

Note
You can navigate to **Objects** > **Intrusion Rules** and click the **Snort 3 All Rules** tab and traverse through all the intrusion rule groups. The parent rule group lists the associated child groups and rule count.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

Rule Group Reporting

The rule groups are reflected in the intrusion events generated and MITRE tactics and techniques are also called out. There are columns for MITRE tactics and techniques and for non-MITRE rule groups for intrusion events. To access the intrusion events, in Firewall Management Center, go to **Analysis** > **Intrusions** > **Events**, and click the **Table View of Events** tab. You can also view the intrusion event fields in the **Unified Events** viewer. In the **Analysis** tab, click **Unified Events**.

In the **Intrusion Events** page, the following fields are added for rule group reporting. Note that you must explicitly enable the mentioned columns.

- MITRE ATT&CK
- Rule Group

For information about these fields, see the section *Intrusion Event Fields* in the *Cisco Secure Firewall Management Center Administration Guide, 7.3*.

Rule Action Logging

From Firewall Management Center 7.2.0 onwards, in the **Intrusion Events** page, the event in the **Inline Result** column displays the same name as the IPS action applied to the rule, so that you can see the action that was applied on the traffic matching the rule.

For the IPS actions, the following table shows the events that are displayed in the **Inline Result** column of the **Intrusion Events** page and **Action** column for **Intrusion Event Type** in the **Unified Events** page.

IPS Action for Snort 3	Inline Result - Firewall Management Center 7.1.0 and earlier	Inline Result -Firewall Management Center 7.2.0 onwards
Alert	Pass	Alert
Block	Dropped/Would Have Dropped/Partially Dropped	Block/Would Block/Partial Block
Drop	Dropped/Would have dropped	Drop/Would drop
Reject	Dropped/Would have dropped	Reject/Would reject

IPS Action for Snort 3	Inline Result - Firewall Management Center 7.1.0 and earlier	Inline Result -Firewall Management Center 7.2.0 onwards
Rewrite	Allow	Rewrite

Important

- In case of a rule without the “Replace” option, the **Rewrite** action is displayed as **Would Rewrite**.
- The **Rewrite** action would also be displayed as **Would Rewrite** if the "Replace" option is specified, but the IPS policy is in Detection mode or the device is in Inline-TAP/Passive mode.

Note In case of backward compatibility (Firewall Management Center 7.2.0 managing a Firewall Threat Defense 7.1.0 device), the events mentioned are applicable only to the Alert IPS action where **Pass** is displayed as **Alert** for events. For all the other actions, the events for Firewall Management Center 7.1.0 are applicable.

Change the Base Policy of an Intrusion Policy

You can choose a different system-provided or custom policy as your base policy.

You can chain up to five custom policies, with four of the five using one of the other four previously created policies as its base policy; the fifth must use a system-provided policy as its base.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Click **Edit** (✎) next to the intrusion policy you want to configure.

Step 3 Choose a policy from the **Base Policy** drop-down list.

Step 4 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

View Snort 2 and Snort 3 Base Policy Mapping

Note Snort 2 is not supported on threat defense Version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the Firewall Management Center guide that matches your Firewall Threat Defense version.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Ensure the **Intrusion Policies** tab is selected.

Step 3 Click **IPS Mapping**.

Step 4 In the **IPS Policy Mapping** dialog box, click **View Mappings** to view the Snort 3 to Snort 2 intrusion policy mapping.

Step 5 Click **OK**.

Synchronize Snort 2 Rules with Snort 3

To ensure that the Snort 2 version settings and custom rules are retained and carried over to Snort 3, the Firewall Management Center provides the synchronization functionality. Synchronization helps Snort 2 rule override settings and custom rules, which you may have altered and added over the last few months or years, to be replicated on the Snort 3 version. This utility helps to synchronize Snort 2 version policy configuration with Snort 3 version to start with similar coverage.

Note Snort 2 is not supported on threat defense Version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the Firewall Management Center guide that matches your Firewall Threat Defense version.

If the Firewall Management Center is upgraded from 6.7 or earlier to 7.0 or later version, the system synchronizes the configuration. If the Firewall Management Center is a fresh 7.0 or later version, you can upgrade to a higher version, and the system will not synchronize any content during upgrade.

Before upgrading a device to Snort 3, if changes are made in Snort 2 version, you can use this utility to have the latest synchronization from Snort 2 version to Snort 3 version so that you start with a similar coverage.

Note On moving to Snort 3, it is recommended that you manage the Snort 3 version of the policy independently and do not use this utility as a regular operation.

Important

- Only the Snort 2 rule overrides and custom rules are copied to Snort 3 and not the other way around. You may not find a one-to-one mapping of all the intrusion rules in Snort 2 and Snort 3. Your changes to rule actions for rules that exist in both versions are synchronized when you perform the following procedure.
- Synchronization *does not* migrate the threshold and suppression settings of any custom or system-provided rules from Snort 2 to Snort 3.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Ensure the **Intrusion Policies** tab is selected.

Step 3 Click **Show Snort 3 Sync status**.

Step 4 Identify the intrusion policy that is out-of-sync.

Step 5 Click the **Sync** icon **Snort out-of-Sync** ( ).

Note
If the Snort 2 and the Snort 3 versions of the intrusion policy are synchronized, then the **Sync** icon is in blue**Snort in-Sync** ( ).

Step 6 Read through the summary and download a copy of the summary if required.

Step 7 Click **Re-Sync**.

Note

- The synchronized settings will be applicable on the Snort 3 intrusion engine only if it is applied on a device, and after a successful deployment.
- Snort 2 custom rules can be converted to Snort 3 using the system-provided tool. If you have any Snort 2 custom rules click the Custom Rules tab and follow the on-screen instructions to convert the rules. For more information, see Convert Snort 2 Custom Rules of a Single Intrusion Policy to Snort 3, on page 2053.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

Manage Intrusion Policies

On the Intrusion Policy page (**Policies** > **Intrusion**) you can view your current custom intrusion policies, along with the following information:

- Number of access control policies and devices are using the intrusion policy to inspect traffic
- In a multidomain deployment, the domain where the policy was created

Note Snort 2 is not supported on threat defense Version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the Firewall Management Center guide that matches your Firewall Threat Defense version.

In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Manage your intrusion policy:

- Create — Click **Create Policy**; see Create a Custom Snort 3 Intrusion Policy , on page 2060.
- Delete — Click **Delete** ( ) next to the policy you want to delete. The system prompts you to confirm and informs you if another user has unsaved changes in the policy. Click **OK** to confirm.

  If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Edit intrusion policy details — Click **Edit** ( ) next to the policy you want to edit. You can edit the **Name**, **Inspection Mode**, and the **Base Policy** of the intrusion policy.
- Edit intrusion policy settings — Click **Snort 3 Version**; see Edit Snort 3 Intrusion Policies, on page 2061.
- Export — If you want to export an intrusion policy to import on another Firewall Management Center, click Export; see the *Exporting Configurations* topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide*.
- Deploy — Choose **Deploy** > **Deployment**; see Deploy Configuration Changes, on page 2055.
- Report — Click **Report**; see the *Generating Current Policy Reports* topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide*. Generates wo reports, one for each policy version.

Access Control Rule Configuration to Perform Intrusion Prevention

An access control policy can have multiple access control rules associated with intrusion policies. You can configure intrusion inspection for any Allow or Interactive Block access control rule, which permits you to match different intrusion inspection profiles against different types of traffic on your network before it reaches its final destination.

Whenever the system uses an intrusion policy to evaluate traffic, it uses an associated *variable set*. Variables in a set represent values commonly used in intrusion rules to identify source and destination IP addresses and ports. You can also use variables in intrusion policies to represent IP addresses in rule suppressions and dynamic rule states.

Tip Even if you use system-provided intrusion policies, Cisco **strongly** recommends you configure the system's intrusion variables to accurately reflect your network environment. At a minimum, modify default variables in the default set.

Understanding System-Provided and Custom Intrusion Policies

Cisco delivers several intrusion policies with the system. By using system-provided intrusion policies, you can take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos sets intrusion and preprocessor rule states, as well as provides the initial configurations for advanced settings. You can use system-provided policies as-is, or you can use them as the base for custom policies. Building custom policies can improve the performance of the system in your environment and provide a focused view of the malicious traffic and policy violations occurring on your network.

Connection and Intrusion Event Logging

When an intrusion policy invoked by an access control rule detects an intrusion and generates an intrusion event, it saves that event to the Management Center. The system also automatically logs the end of the connection where the intrusion occurred to the Management Center database, regardless of the logging configuration of the access control rule.

Access Control Rule Configuration and Intrusion Policies

The number of unique intrusion policies you can use in a single access control policy depends on the model of the target devices; more powerful devices can handle more. Every unique **pair** of intrusion policy and variable set counts as one policy. Although you can associate a different intrusion policy-variable set pair with each Allow and Interactive Block rule (as well as with the default action), you cannot deploy an access control policy if the target devices have insufficient resources to perform inspection as configured.

Configure an Access Control Rule to Perform Intrusion Prevention

You must be an Admin, Access Admin, or Network Admin to perform this task.

Procedure

Step 1 In the access control policy editor, create a new rule or edit an existing rule; see the *Access Control Rule Components* topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide*.

Step 2 Ensure the rule action is set to **Allow**, **Interactive Block**, or **Interactive Block with reset**.

Step 3 Click **Inspection**.

Step 4 Choose a system-provided or a custom intrusion policy, or choose **None** to disable intrusion inspection for traffic that matches the access control rule.

Step 5 If you want to change the variable set associated with the intrusion policy, choose a value from the **Variable Set** drop-down list.

Step 6 Click **Save** to save the rule.

Step 7 Click **Save** to save the policy.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

Deploy Configuration Changes

After you change configurations, deploy them to the affected devices.

Note This topic covers the basic steps involved in deploying configuration changes. We *strongly* recommend that you refer the *Deploy Configuration Changes* topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide* to understand the prerequisites and implications of deploying the changes before proceeding with the steps.

Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic.

Procedure

Step 1 On the Secure Firewall Management Center menu bar, click **Deploy** and choose **Deployment**.

The GUI page lists the devices with out-of-date configurations having **Pending**status.

- The **Modified By** column lists the users who have modified the policies or objects. Expand the device listing to view the users who have modified the policies for each policy listing.

  Note
  Usernames are not provided for deleted policies and objects.

- The **Inspect Interruption** column indicates if traffic inspection interruption might occur in the device during deployment.

  If this column is blank for a device, it indicates that there will be no traffic inspection interruptions on that device during deployment.

- The **Last Modified Time** column specifies the last time you made configuration changes.
- The **Preview** column allows you to preview the changes for the next deployment.
- The **Status** column provides the status for each deployment.

Step 2 Identify and choose the devices on which you want to deploy configuration changes.

- Search—Search for the device name, type, domain, group, or status in the search box.
- Expand—Click **Expand Arrow** (>) to view device-specific configuration changes to be deployed.

  When you check a check box adjacent to a device, all the changes made to the device and listed under the device, are pushed for deployment. However, you can use **Policy selection** ( ) to select individual policies or specific configurations to deploy while withholding the remaining changes without deploying them.

  Note
  - When the status in the **Inspect Interruption** column indicates (**Yes**) that deploying will interrupt inspection, and perhaps traffic, on a Firewall Threat Defense device, the expanded list indicates the specific configurations causing the interruption with the **Inspect Interruption** ( ).
  - When there are changes to interface groups, security zones, or objects, the impacted devices are shown as out-of-date on the Firewall Management Center. To ensure that these changes take effect, the policies with these interface groups, security zones, or objects, also need to be deployed along with these changes. The impacted policies are shown as out-of-date on the **Preview**page on the Firewall Management Center.

Step 3 Click **Deploy**.

Step 4 If the system identifies errors or warnings in the changes to be deployed, it displays them in the **Validation Messages** window. To view complete details, click the arrow icon before the warnings or errors.

You have the following choices:

- Deploy—Continue deploying without resolving warning conditions. You cannot proceed if the system identifies errors.
- Close—Exit without deploying. Resolve the error and warning conditions, and attempt to deploy the configuration again.

What to do next

During deployment, if there is a deployment failure, there is a possibility that the failure may impact traffic. However, it depends on certain conditions. If there are specific configuration changes in the deployment, the deployment failure may lead to traffic being interrupted. For details, see the Deploy Configuration Changes topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide*.

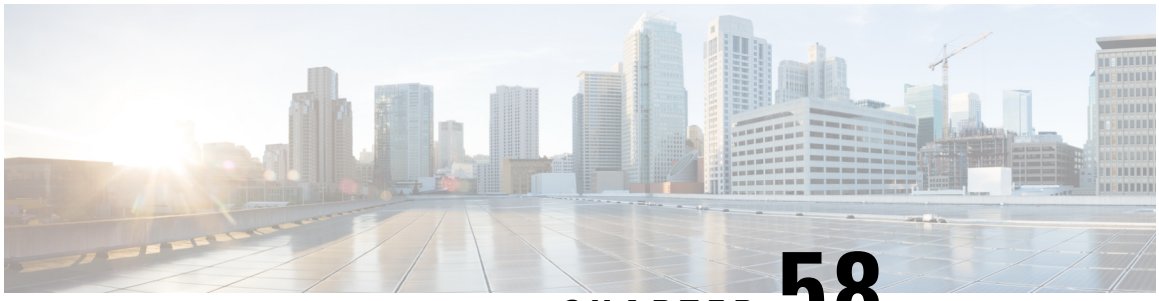

CHAPTER 58

Tune Intrusion Policies Using Rules

This chapter provides information on custom rules in Snort 3, intrusion rule action, intrusion event notification filters in an intrusion policy, converting Snort 2 custom rules to Snort 3, and adding rule groups with custom rules to an intrusion policy.

- Overview of Tuning Intrusion Rules, on page 2073
- Intrusion Rule Types, on page 2074
- Prerequisites for Network Analysis and Intrusion Policies, on page 2075
- Custom Rules in Snort 3, on page 2075
- View Snort 3 Intrusion Rules in an Intrusion Policy, on page 2078
- Intrusion Rule Action, on page 2079
- Intrusion Event Notification Filters in an Intrusion Policy, on page 2080
- Add Intrusion Rule Comments, on page 2085
- Snort 2 Custom Rules Conversion to Snort 3, on page 2085
- Add Custom Rules to Rule Groups, on page 2087
- Add Rule Groups with Custom Rules to an Intrusion Policy, on page 2088
- Manage Custom Rules in Snort 3, on page 2089
- Delete Custom Rules, on page 2090
- Delete Rule Groups, on page 2090

Overview of Tuning Intrusion Rules

You can configure rule states and other settings for shared object rules, standard text rules, and inspector rules.

You enable a rule by setting its rule state to Alert or to Block. Enabling a rule causes the system to generate events on traffic matching the rule. Disabling a rule stops processing of the rule. You can also set your intrusion policy so that a rule set to Block generates events on, and drops, matching traffic.

You can filter rules to display a subset of rules, enabling you to select the exact set of rules where you want to change rule states or rule settings.

When an intrusion rule or rule argument requires a disabled inspector, the system automatically uses it with its current configuration even though it remains disabled in the network analysis policy's web interface.

Note We recommend that you do not modify shared object rules and you only enable or disable these rules for your threat defense device. To create custom Snort rules, contact Cisco support.

Intrusion Rule Types

An intrusion rule is a specified set of keywords and arguments that the system uses to detect attempts to exploit vulnerabilities in your network. As the system analyzes network traffic, it compares packets against the conditions specified in each rule, and triggers the rule if the data packet meets all the conditions specified in the rule.

An intrusion policy contains:

- *intrusion rules*, which are subdivided into *shared object rules* and *standard text rules*
- *inspector rules*, which are associated with a detection option of the packet decoder or with one of the inspectors included with the system

The following table summarizes attributes of these rule types:

Table 109: Intrusion Rule Types

Type	Generator ID (GID)	Snort ID (SID)	Source	Can Copy?	Can Edit?
shared object rule	3	lower than 1000000	Cisco Talos Intelligence Group (Talos)	yes	limited
standard text rule	1 (Global domain or legacy GID)	lower than 1000000	Talos	yes	limited
	1000 - 2000 (descendant domain)	1000000 or higher	Created or imported by user	yes	yes
preprocessor rule	decoder- or preprocessor-specific	lower than 1000000	Talos	no	no
		1000000 or higher	Generated by the system during option configuration	no	no

You cannot save changes to any rule created by Talos, but you can save a copy of a modified rule as a custom rule. You can modify either variables used in the rule or rule header information (such as source and destination ports and IP addresses). In a multidomain deployment, rules created by Talos belong to the Global domain. Administrators in descendant domains can save local copies of the rules, which they can then edit.

For the rules it creates, Talos assigns default rule states in each default intrusion policy. Most preprocessor rules are disabled by default and must be enabled if you want the system to generate events for preprocessor rules and, in an inline deployment, drop offending packets.

Prerequisites for Network Analysis and Intrusion Policies

To allow the Snort inspection engine to process traffic for intrusion and malware analysis, you must have the IPS license enabled for the Firewall Threat Defense device.

You must be an Admin user to manage network analysis, intrusion policies, and perform migration tasks.

Custom Rules in Snort 3

You can create a custom intrusion rule by importing a local rule file. The rule file can either have a `.txt` or `.rules` extension. The system saves the custom rule in the local rule category, regardless of the method you used to create it. A custom rule must belong to a rule group. However, a custom rule can be a part of two or more groups as well.

When you create a custom intrusion rule, the system assigns it a unique rule number, which has the format `GID:SID:Rev`. The elements of this number are:

- **GID**—Generator ID. For custom rules, it is not necessary to specify the GID. The system automatically generates the GID based on whether you are in the Global domain or a sub-domain while uploading the rules. For all standard text rules, this value is 2000 for a Global domain.
- **SID**—Snort ID. Indicates whether the rule is a local rule of a system rule. When you create a new rule, assign a unique SID to the rule.

  SID numbers for local rules start at 1000000, and the SID for each new local rule is incremented by one.
- **Rev**—The revision number. For a new rule, the revision number is one. Each time you modify a custom rule the revision number should be incremented by one.

In a custom standard text rule, you set the rule header settings and the rule keywords and arguments. You can use the rule header settings to focus the rule to only match traffic using a specific protocol and traveling to or from specific IP addresses or ports.

To check whether a SID is enabled or disabled, verify the entries in the snort.lua file located in the `./file-contents/ngfw/var/sf/detection_engines/<id>/ips/<id>` directory.

- If the SID is disabled by default, no entry will be present in the file.
- If the SID is manually enabled, you will see an entry with **enable:yes**.
- If the SID is disabled after being manually enabled, the entry remains in the file and will display **enable:no**.

Note

- Snort 3 custom rules cannot be edited. Ensure that custom rules have a valid classification message for `classtype` within the rule text. If you import a rule without a classification or wrong classification, then delete and recreate the rule.
- You can create custom intrusion rules using Snort 3. However, support for tuning and troubleshooting these rules is not available currently.
- The *classtype* in a Snort rule assigns a classification to the rule indicating the type of attack that is associated with an event. A priority level of 1-4 is also associated with each *classtype*. However, the priority level for certain *classtypes* on the Threat Defense device do not match the open-source Snort *classtype* priority levels that are mentioned in the Snort documentation. For example, *tcp-connection* has a priority of 4 in open-source Snort while a priority of 3 is assigned to it on the Threat Defense device.

Sensitive Data Detection in Snort 3

Sensitive data such as social security numbers, credit card numbers, emails, and so on may be leaked onto the internet, intentionally or accidentally. Sensitive data detection is used to detect and generate events on possible sensitive data leakage. Events are generated only if there is a transfer of significant amount of Personally Identifiable Information (PII) data. Sensitive data detection can mask PII in the output of events.

sd_pattern Option

Use the `sd_pattern` IPS option to detect and filter PII. This information includes credit card numbers, U.S. Social Security numbers, phone numbers, and email addresses. A regular expression (regex) syntax is available for defining your own PII.

The sd_pattern option has the following settings:

- Pattern—An implicit, required setting that specifies the regular expression to look for in the PDU. The regex must be written in PCRE syntax.
- Threshold—An explicit, optional setting that specifies the number of matches in the PDU required to generate an event.

  The `sd_pattern` as IPS rule option is available in Snort with no requirements for additional inspectors. The rule option's syntax is:

  ```
  sd_pattern: "<pattern>"[, threshold <count>];
  ```

  For example:

  ```
  sd_pattern:"credit_card", threshold 2;
  ```

Built-in Patterns

There are five built-in patterns for sensitive data. To use the built-in patterns in the "pattern" setting, you must specify the name of the PII type that needs to be matched and the necessary regex is substituted for it. The PII name and regex mappings or patterns are described as follows:

- credit_card—

  ```
  \d{4}\D?\d{4}\D?\d{2}\D?\d{2}\D?\d{3,4}
  ```

- us_social—

```
[0-8]\d{2}-\d{2}-\d{4}
```

- us_social_nodashes—

```
[0-8]\d{8}
```

- Email—

```
[a-zA-Z0-9!#$%&'*+\/=?^_`{|}~-]+(?:\.[a-zA-Z0-9!#$%&'*+\/=?^_`{|}~-]+)*@(?:[a-zA-Z0-9]
(?:[a-zA-Z0-9-]*[a-zA-Z0-9])?\.)+[a-zA-Z0-9](?:[a-zA-Z0-9-]*[a-zA-Z0-9])?
```

- us_phone—

```
(?:\+?1[-\.\s]?)?\(?([2-9][0-8]\d)\)?[-\.\s]([2-9]\d{2})[-\.\s](\d{4})
```

PII Name	Pattern
credit_card	`\d{4}\D?\d{4}\D?\d{2}\D?\d{2}\D?\d{3,4}`
us_social	`[0-8]\d{2}-\d{2}-\d{4}`
us_social_nodashes	`[0-8]\d{8}`
Email	`[a-zA-Z0-9!#$%&'*+\/=?^_`{\|}~-]+(?:\.[a-zA-Z0-9!#$%&'*+\/=?^_`{|}~-]+)*@(?:[a-zA-Z0-9](?:[a-zA-Z0-9-]*[a-zA-Z0-9])?\.)+[a-zA-Z0-9](?:[a-zA-Z0-9-]*[a-zA-Z0-9])?`
us_phone	`(?:\+?1[-\.\s]?)?\(?([2-9][0-8]\d)\)?[-\.\s]([2-9]\d{2})[-\.\s](\d{4})`

Masking for data matching these patterns only work with system-provided rules or built-in patterns for Credit Cards, U.S. Social Security numbers, emails, and U.S. phone numbers. Masking does not work for custom rules or user-defined PII patterns. Rules are available in the Lightweight Security Package (LSP) for sensitive data, gid:13. By default, they are not enabled in any system-provided policy.

The sensitive data rules in LSP cover all built-in patterns and have the following threshold values:

- credit_card: 2
- us_social: 2
- us_social_nodashes: 20
- email: 20
- us_phone: 20

You can use the sd_pattern option to create custom rules and modify existing rules. To do this, use the Snort 3 intrusion policy interface.

An example of a rule with sd_pattern with a custom pattern and threshold:

alert tcp (sid: 1000000001; sd_pattern:"[\w-\.]+@([\w-]+\.)+[\w-]{2,4}",threshold 4; msg: "email, threshold 4")

Examples

An example of custom rules using sensitive data detection:

Rule with built-in pattern:

```
alert tcp (
        msg:"SENSITIVE-DATA Email";
```

```
        flow:only_stream;
        pkt_data;
        sd_pattern:"email", threshold 5;
        service:http, smtp, ftp-data, imap, pop3;
        gid:2000;
        sid:1000001;
)
```

Rule with custom pattern

```
alert tcp (
        msg:"SENSITIVE-DATA US phone numbers";
        flow:only_stream;
        file_data;
        sd_pattern:"+?3?8?(0[\s\.-]\d{2}[\s\.-]\d{3}[\s\.-]\d{2}[\s\.-]\d{2})", threshold
2;
        service:http, smtp, ftp-data, imap, pop3;
        gid:2000;
        sid:1000002;
)
```

Here are some more examples of complete Snort IPS rules with built-in sensitive data patterns:

- alert tcp (sid:1; msg:"Credit Card"; sd_pattern:"credit_card", threshold 2;)
- alert tcp (sid:2; msg:"US Social Number"; sd_pattern:"us_social", threshold 2;)
- alert tcp (sid:3; msg:"US Social Number No Dashes"; sd_pattern:"us_social_nodashes", threshold 2;)
- alert tcp (sid:4; msg:"US Phone Number"; sd_pattern:"us_phone", threshold 2;)
- alert tcp (sid:5; msg:"Email"; sd_pattern:"email", threshold 2;)

Disabling data masking is not supported in the Secure Firewall Management Center and Secure Firewall Device Manager.

For information how to add custom rules to rule groups, see Add Custom Rules to Rule Groups, on page 2087.

View Snort 3 Intrusion Rules in an Intrusion Policy

You can adjust how rules are displayed in the intrusion policy. You can also display the details for a specific rule to see rule settings, rule documentation, and other rule specifics.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Click **Snort 3 Version** next to the policy.

Step 3 While viewing the rules, you can:

- Filter the rules.
- Choose a rule group to see rules related to that group.
- View an intrusion rule's details.
- View rule comments.
- View rule documentation.

See Edit Snort 3 Intrusion Policies, on page 2061 for details on performing these tasks.

Intrusion Rule Action

Intrusion rule action allows you to enable or disable the rule within an individual intrusion policy, as well as specify which action the system takes if monitored conditions trigger the rule.

The Cisco Talos Intelligence Group (Talos) sets the default action of each intrusion and inspector rule in each default policy. For example, a rule may be enabled in the Security over Connectivity default policy and disabled in the Connectivity over Security default policy. Talos sometimes uses a rule update to change the default action of one or more rules in a default policy. If you allow rule updates to update your base policy, you also allow the rule update to change the default action of a rule in your policy when the default action changes in the default policy you used to create your policy (or in the default policy it is based on). Note, however, that if you have changed the rule action, the rule update does not override your change.

When you create an intrusion rule, it inherits the default actions of the rules in the default policy you use to create your policy.

Intrusion Rule Action Options

In an intrusion policy, you can set a rule's action to the following values:

Alert

You want the system to detect a specific intrusion attempt and generate an intrusion event when it finds matching traffic. When a malicious packet crosses your network and triggers the rule, the packet is sent to its destination and the system generates an intrusion event. The malicious packet reaches its target, but you are notified through the event logging.

Block

You want the system to detect a specific intrusion attempt, drop the packet containing the attack, and generate an intrusion event when it finds matching traffic. The malicious packet never reaches its target, and you are notified through the event logging.

Disable

You do not want the system to evaluate matching traffic.

Note Choosing either the **Alert** or **Block** options enables the rule. Choosing **Disable** disables the rule.

We **strongly** recommend that you **do not** enable all the intrusion rules in an intrusion policy. The performance of your managed device is likely to degrade if all rules are enabled. Instead, tune your rule set to match your network environment as closely as possible.

Set Intrusion Rule Action

Intrusion rule actions are policy-specific.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Click **Snort 3 Version** next to the policy you want to edit.

Tip
This page shows the total number of:

- disabled rules
- enabled rules set to Alert
- enabled rules set to Block
- overridden rules

Step 3 Choose the rule or rules where you want to set the rule action.

Step 4 Choose one of the rule actions from the **Rule Action** drop-down list. See Edit Snort 3 Intrusion Policies, on page 2061 for more information about the different rule actions.

Step 5 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

Intrusion Event Notification Filters in an Intrusion Policy

The importance of an intrusion event can be based on frequency of occurrence, or on source or destination IP address. In some cases you may not care about an event until it has occurred a certain number of times. For example, you may not be concerned if someone attempts to log into a server until they fail a certain number of times. In other cases, you may only need to see a few occurrences to know there is a widespread problem. For example, if a DoS attack is launched against your web server, you may only need to see a few occurrences of an intrusion event to know that you need to address the situation. Seeing hundreds of the same event only overwhelms your system.

Intrusion Event Thresholds

You can set thresholds for individual rules to limit the number of times the system logs and displays an intrusion event based on how many times the event is generated within a specified time period. This can prevent you from being overwhelmed with a large number of identical events. You can set thresholds per shared object rule, standard text rule, or inspector rule.

Set Intrusion Event Thresholds

To set a threshold, first specify the thresholding type.

Table 110: Thresholding Options

Option	Description
Limit	Logs and displays events for the specified number of packets (specified by the Count argument) that trigger the rule during the specified time period. For example, if you set the type to **Limit**, the **Count** to `10`, and the **Seconds** to `60`, and 14 packets trigger the rule, the system stops logging events for the rule after displaying the first 10 that occur within the same minute.
Threshold	Logs and displays a single event when the specified number of packets (specified by the Count argument) trigger the rule during the specified time period. Note that the counter for the time restarts after you hit the threshold count of events and the system logs that event. For example, you set the type to **Threshold**, **Count** to `10`, and **Seconds** to `60`, and the rule triggers 10 times by second 33. The system generates one event, then resets the Seconds and Count counters to 0. The rule then triggers another 10 times in the next 25 seconds. Because the counters reset to 0 at second 33, the system logs another event.
Both	Logs and displays an event once per specified time period, after the specified number (count) of packets trigger the rule. For example, if you set the type to **Both**, **Count** to two, and **Seconds** to `10`, the following event counts result: • If the rule is triggered once in 10 seconds, the system does not generate any events (the threshold is not met) • If the rule is triggered twice in 10 seconds, the system generates one event (the threshold is met when the rule triggers the second time) • If the rule is triggered four times in 10 seconds, the system generates one event (the threshold is met when the rule triggers the second time, and following events are ignored)

Secondly, specify tracking, which determines whether the event threshold is calculated per source or destination IP address.

Table 111: Thresholding IP Options

Option	Description
Source	Calculates event instance count per source IP address.
Destination	Calculates event instance count per destination IP address.

Finally, specify the number of instances and time period that define the threshold.

Table 112: Thresholding Instance/Time Options

Option	Description
Count	The number of event instances per specified time period per tracking IP address required to meet the threshold.

Option	Description
Seconds	The number of seconds that elapse before the count resets. If you set the threshold type to **limit**, the tracking to **Source IP**, the **count** to `10`, and the **seconds** to `10`, the system logs and displays the first 10 events that occur in 10 seconds from a given source port. If only 7 events occur in the first 10 seconds, the system logs and displays those; if 40 events occur in the first 10 seconds, the system logs and displays 10, then begins counting again when the 10-second time period elapses.

Note that you can use intrusion event thresholding alone or in any combination with rate-based attack prevention, the `detection_filter` keyword, and intrusion event suppression.

Tip You can also add thresholds from within the packet view of an intrusion event.

Set Threshold for an Intrusion Rule in Snort 3

You can set a single threshold for a rule from the Rule Detail page. Adding a threshold overwrites any existing threshold for the rule. The threshold you set for an intrusion rule is applied to each packet thread. However, the configuration is fully applied only within the context of a unique flow. There may be more alerts on different network flows, but there will not be fewer alerts than the configured number.

Procedure

Step 1 Choose **Objects** > **Intrusion Rules** .

Step 2 Click **Snort 3 All Rules** tab.

Step 3 From an intrusion rule's Alert Configuration column, click the **None** link.

Step 4 Click **Edit** ( ).

Step 5 In the Alert Configuration window, click the **Threshold** tab.

Step 6 From the **Type** drop-down list, choose the type of threshold you want to set:

- Choose **Limit** to limit notification to the specified number of event instances per time period.
- Choose **Threshold** to provide notification for each specified number of event instances per time period.
- Choose **Both** to provide notification once per time period after a specified number of event instances.

Step 7 Choose **Source** or **Destination** in the **Track By** field to indicate whether you want the event instances tracked by source or destination IP address.

Step 8 Enter the number of event instances you want to use as your threshold in the **Count** field.

Step 9 Enter a number that specifies the time period, in seconds, for which event instances are tracked in the **Seconds** field.

Step 10 Click **Save**.

Refer to the video Snort 3 Suppression and Threshold for additional support and information.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 259.

View and Delete Intrusion Event Thresholds

To view or delete an existing threshold setting for a rule, use the Rules Details view to display the configured settings for a threshold and see if they are appropriate for your system. If they are not, you can add a new threshold to overwrite the existing values.

Procedure

Step 1 Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Choose the rule with a configured threshold as shown in the **Alert Configuration** column (the **Alert Configuration** column displays **Threshold** as a link for the rule).

Step 4 To remove the threshold for the rule, click **Threshold** link in the **Alert Configuration** column.

Step 5 Click **Edit** ( ).

Step 6 Click **Threshold** tab.

Step 7 Click **Reset**.

Step 8 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

Intrusion Policy Suppression Configuration

You can suppress intrusion event notification when a specific IP address or range of IP addresses triggers a specific rule or inspector. This is useful for eliminating false positives. For example, if you have a mail server that transmits packets that look like a specific exploit, you might suppress event notification for that event when it is triggered by your mail server. The rule triggers for all packets, but you only see events for legitimate attacks.

Intrusion Policy Suppression Types

Note that you can use intrusion event suppression alone or in any combination with rate-based attack prevention, the `detection_filter` keyword, and intrusion event thresholding.

Tip You can add suppressions from within the packet view of an intrusion event. You can also access suppression settings by using the **Alert Configuration** column on the intrusion rules editor page (**Objects** > **Intrusion Rules** > **Snort 3 All Rules**).

Set Suppression for an Intrusion Rule in Snort 3

You can set one or more suppressions for a rule in your intrusion policy.

Before you begin

Ensure you create the required network objects to be added for source or destination suppression.

Procedure

Step 1 Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Click the **None** link in the intrusion rule's Alert Configuration column,.

Step 4 Click **Edit** ( ).

Step 5 From the **Suppressions** tab, click the add icon **Add** ( ) next to any of the following options:

- Choose **Source Networks** to suppress events generated by packets originating from a specified source IP address.
- Choose **Destination Networks** to suppress events generated by packets going to a specified destination IP address.

Step 6 Select any of the preset networks in the **Network** drop-down list.

Step 7 Click **Save**.

Step 8 (Optional) Repeat the last three steps if required.

Step 9 Click **Save** in the Alert Configuration window.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

View and Delete Suppression Conditions

You may want to view or delete an existing suppression condition. For example, you can suppress event notification for packets originating from a mail server IP address because the mail server normally transmits packets that look like exploits. If you then decommission that mail server and reassign the IP address to another host, you should delete the suppression conditions for that source IP address.

Procedure

Step 1 Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Choose the rule for which you want to view or delete suppressions.

Step 4 Click **Suppression** in the **Alert Configuration** column.

Step 5 Click **Edit** ( ).

Step 6 Click **Suppressions** tab.

Step 7 Remove the suppression by clicking **Clear** ( ) next to the suppression.

Step 8 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

Add Intrusion Rule Comments

You can add comments to rules in your intrusion policy. Comments added this way are policy-specific; that is, comments you add to a rule in one intrusion policy are not visible in other intrusion policies.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Click **Snort 3 Version** next to the policy you want to edit.

Step 3 In the right side of the page where all the rules are listed, choose the rule where you want to add a comment.

Step 4 Click **Comment** ( ) under the **Comments** column.

Step 5 In the **Comments** field, enter the rule comment.

Step 6 Click **Add Comment**.

Step 7 Click **Save**.

Tip
The system displays a **Comment** ( ) next to the rule in the Comments column.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

Snort 2 Custom Rules Conversion to Snort 3

If you are using custom rules, make sure you are prepared to manage that rule set for Snort 3 prior to conversion from Snort 2 to Snort 3. If you are using a rule set from a third-party vendor, contact that vendor to confirm that their rules will successfully convert to Snort 3 or to obtain a replacement rule set written natively for Snort 3. If you have custom rules that you have written yourself, familiarize with writing Snort 3 rules prior to conversion, so you can update your rules to optimize Snort 3 detection after conversion. See the links below to learn more about writing rules in Snort 3.

- https://blog.snort.org/2020/08/how-rules-are-improving-in-snort-3.html

- https://blog.snort.org/2020/10/talos-transition-to-snort-3.html

You can refer to other blogs at https://blog.snort.org/ to learn more about Snort 3 rules.

Important Snort 2 network analysis policy (NAP) settings *cannot* be copied to Snort 3 automatically. NAP settings have to be manually replicated in Snort 3.

Note Snort 2 is not supported on threat defense Version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the Firewall Management Center guide that matches your Firewall Threat Defense version.

Convert all Snort 2 Custom Rules across all Intrusion Policies to Snort 3

Procedure

Step 1 Choose **Objects** > **Intrusion Rules** .

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Ensure **All Rules** is selected in the left pane.

Step 4 Click the **Tasks** drop-down list and choose:

- **Convert Snort 2 rules and import**—To automatically convert all the Snort 2 custom rules across all the intrusion policies to Snort 3 and import them into Firewall Management Center as Snort 3 custom rules.
- **Convert Snort 2 rules and download**—To automatically convert all the Snort 2 custom rules across all the intrusion policies to Snort 3 and download them into your local system.

Step 5 Click **OK**.

Note

- If you selected **Convert and import** in the previous step, then all the converted rules are saved under a newly created rule group **All Snort 2 Converted Global** under **Local Rules**.
- If you selected **Convert and download** in the previous step, then save the rules file locally. You can review the converted rules in the downloaded file and later upload them by following the steps in Add Custom Rules to Rule Groups, on page 2087.

Refer to the video Converting Snort 2 Rules to Snort 3 for additional support and information.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

Convert Snort 2 Custom Rules of a Single Intrusion Policy to Snort 3

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 In the **Intrusion Policies** tab, click **Show Snort 3 Sync status**.

Step 3 Click the **Sync** icon **Snort out-of-Sync** (🔶) of the intrusion policy.

Note
If the Snort 2 and the Snort 3 versions of the intrusion policy are synchronized, then the **Sync** icon is in blue**Snort in-Sync** (🔷). It indicates that there are no custom rules to be converted.

Step 4 Read through the summary and click the **Custom Rules** tab.

Step 5 Choose:

- **Import converted rules to this policy**—To convert the Snort 2 custom rules in the intrusion policy to Snort 3 and import them into Firewall Management Center as Snort 3 custom rules.
- **Download converted rules**—To convert the Snort 2 custom rules in the intrusion policy to Snort 3 and download them into your local system. You can review the converted rules in the downloaded file and later upload the file by clicking the upload icon.

Step 6 Click **Re-Sync**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

Add Custom Rules to Rule Groups

Uploading custom rules in the Firewall Management Center adds the custom rules that you have created locally to the list of all the Snort 3 rules.

Procedure

Step 1 Choose .

Policies > Intrusion No conref was created

Objects > **Intrusion Rules**

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Click the **Tasks** drop-down list.

Step 4 Click **Upload Snort 3 Rules**.

Step 5 Drag and drop the `.txt` or `.rules` file that contains the Snort 3 custom rules that you have created.

Step 6 Click **OK**.

Note
If there are any errors in the selected file, then you cannot proceed further. You can download the error file and **Replace File** link to upload version 2 of the file, after fixing the errors.

Step 7 Associate rules to a rule group to add the new rules to that group.

You can also create a new custom rule group (by clicking the **Create New Custom Rule Group** link) and then add the rules to the new group.

Note
If there are no existing local rule groups, then proceed by clicking **Create New Custom Rule Group to proceed**. Enter a **Name** for the new rule group and click **Save**.

Step 8 Choose either of the following:

- **Merge Rules** to merge the new rules that you are adding with the existing rules in the rule group.
- **Replace all rules in the group with file contents** to replace all the exiting rules with the new rules that you are adding.

Note
If you chose more than one rule group in the previous step, then only the **Merge Rules** option is available.

Step 9 Click **Next**.

Review the summary to know the new rule IDs that are being added and optionally download it.

Step 10 Click **Finish**.

Important The rule action of all the uploaded rules is in the disabled state. You have to change them to the required state to ensure the rules are active.

What to do next

- Uploading custom rules in the Firewall Management Center adds the custom rules that you have created to the list of all the Snort 3 rules. To enforce these custom rules on the traffic, add and enable these rules in the required intrusion policies. For information on adding rule groups with custom rules to an intrusion policy, see Add Rule Groups with Custom Rules to an Intrusion Policy, on page 2088. For information on enabling custom rules, see Manage Custom Rules in Snort 3, on page 2089.
- Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

Add Rule Groups with Custom Rules to an Intrusion Policy

Custom rules that are uploaded in the system have to be enabled in an intrusion policy to enforce those rules on the traffic. After uploading custom rules on Firewall Management Center, add the rule group with the new custom rules in the intrusion policy.

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 In the **Intrusion Policies** tab, click the **Snort 3 Version** of the intrusion policy.

Step 3 Click **Add** (+) next to the Rule Groups search bar.

Step 4 In the **Add Rule Groups** window, click the **Expand Arrow** (>) icon next to a rule group to expand the local rule group.

Step 5 Check the check box next to the uploaded custom rules group.

Step 6 Click **Save**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

Manage Custom Rules in Snort 3

Custom rules that are uploaded in the system have to be added to an intrusion policy and enabled to enforce those rules on the traffic. You can enable the uploaded custom rules across all policies or selectively on individual policies.

Follow the steps to enable custom rules in one or many intrusion policies:

Procedure

Step 1 Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Expand **Local Rules**.

Step 4 Select the required rule group.

Step 5 Select the rules by checking the check boxes next to them.

Step 6 Select **Per Intrusion Policy** from the **Rule Actions** drop-down list.

Step 7 Choose:

- **All Policies**—to have the same rule actions for all the rules to be added.
- **Per Intrusion Policy**—to have different rule actions for each intrusion policy.

Step 8 Set the rule actions:

- If you selected All Policies in the previous step, then select the required rule action from the **Select Override state** drop-down list.
- If you selected Per Intrusion Policy in the previous step, then select the **Rule Action** against the policy name. To add more policies, click **Add Another**.

Step 9 Optionally, add a comment in the **Comments** text box.

Step 10 Click **Save**.

What to do next

Deploy the changes on the device. See, Deploy Configuration Changes, on page 2055.

Delete Custom Rules

Procedure

Step 1 Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Expand **Local Rules** in the left pane.

Step 4 Check the check boxes of the rules you want to delete.

Step 5 Ensure that the rule action for all the rules that you select is **Disable**.

If required, follow the steps below to disable the rule action for multiple selected rules:

a) From the **Rule Actions** drop-down box, select **Per Intrusion Policy**.
b) Select **All Policies** radio button.
c) Select **Disable** from the **Select Override state** drop-down list.
d) Click **Save**.
e) Check the check boxes of the rules you want to delete.

Step 6 From the **Rule Actions** drop-down list, select **Delete**.

Step 7 Click **Delete** in the Delete Rules pop-up window.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

Delete Rule Groups

Before you begin

Exclude the rule group you want to delete from all intrusion policies where you have included it. For steps on excluding a rule group from an intrusion policy, see Edit Snort 3 Intrusion Policies, on page 2061.

Procedure

Step 1 Choose **Objects** > **Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Expand **Local Rules** in the left pane.

Step 4 Select the rule group to be deleted.

Step 5 Ensure the rule action for all the rules in the group is set to **Disable** before proceeding.

If the rule action for any of the rules is anything other than **Disable**, then you cannot delete the rule group. If required, follow the steps below to disable the rule action for all the rules:

a) Check the check box below the **Rule Actions** drop-down list to select all the rules in the group.
b) From the **Rule Actions** drop-down box, select **Per Intrusion Policy**.
c) Select **All Policies** radio button.
d) Select **Disable** from the **Select Override state** drop-down list.
e) Click **Save**.

Step 6 Click the **Delete** ( ) next to the rule group.

Step 7 Click **OK** in the Delete Rule Group pop-up window.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

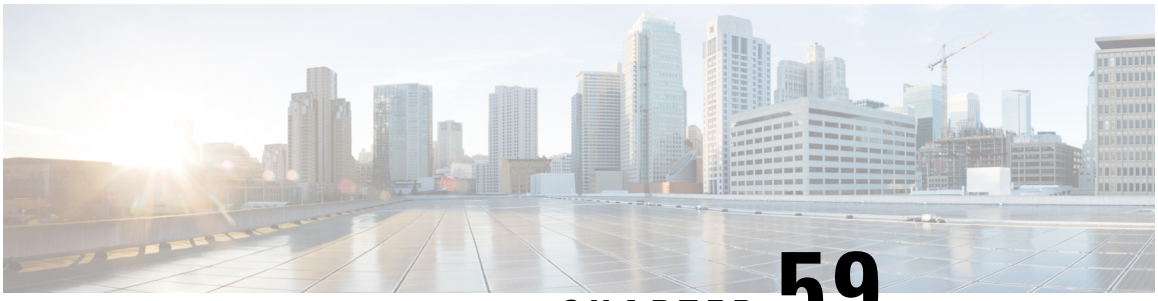

CHAPTER 59

Tailor Intrusion Protection for Your Network Assets

This chapter provides an insight into Secure Firewall recommended rules and generating and applying Secure Firewall recommended rules.

- Snort 3 Rule Changes in LSP Updates , on page 2093
- Overview of Secure Firewall Recommended Rules, on page 2093
- Prerequisites for Network Analysis and Intrusion Policies, on page 2094
- Generate New Secure Firewall Recommendations in Snort 3, on page 2095

Snort 3 Rule Changes in LSP Updates

During regular Snort 3 Lightweight Security Package (LSP) updates, an existing system-defined intrusion rule may be replaced with a new intrusion rule. There could be possibilities of a single rule being replaced with multiple rules, or multiple rules being replaced with a single rule. This occurs when better detection is possible for which rules are combined or expanded. For better management, some existing system-defined rules may also be removed as a part of the LSP update.

To get notifications for changes to any *overridden* system-defined rules during LSP updates, ensure that the **Retain user overrides for deleted Snort 3 rules** check box is checked.

To navigate to the **Retain user overrides for deleted Snort 3 rules** check box, click **System** (⚙), and then choose **Configuration** > **Intrusion Policy Preferences**.

By default this check box is checked. When this check box is checked, the system retains the rule overrides in the new replacement rules that are added as a part of the LSP update. The notifications are shown in the **Tasks** tab under the Notifications icon that is located next to **System** (⚙).

Overview of Secure Firewall Recommended Rules

You can use intrusion rule recommendations to target vulnerabilities associated with host assets detected in the network. For example, operating systems, servers, and client application protocols. This allows you to tailor your intrusion policy to the specific needs of your monitored network.

The system makes an individual set of recommendations for each intrusion policy. It typically recommends rule state changes for standard text rules and shared object rules. However, it can also recommend changes for inspector and decoder rules.

When you generate rule state recommendations, you can use the default settings or configure advanced settings. Advanced settings allow you to:

- Redefine which hosts on your network the system monitors for vulnerabilities
- Influence which rules the system recommends based on rule overhead
- Specify whether to generate recommendations to disable rules

You can also choose to use the recommendations immediately or review the recommendations (and affected rules) before accepting them.

Choosing to use recommended rule states adds a read-only Secure Firewall Recommendations layer to your intrusion policy, and subsequently choosing not to use recommended rule states removes the layer.

You can schedule a task to generate recommendations automatically based on the most recently saved configuration settings in your intrusion policy.

The system does not change rule states that you set manually such as:

- Manually setting the states of specified rules *before* you generate recommendations prevents the system from modifying the states of those rules in the future.
- Manually setting the states of specified rules *after* you generate recommendations overrides the recommended states of those rules.

Tip The intrusion policy report can include a list of rules with rule states that differ from the recommended state.

While displaying the recommendation-filtered Rules page, or after accessing the Rules page directly from the navigation panel or the Policy Information page, you can manually set rule states, sort rules, and take any of the other actions available on the Rules page, such as suppressing rules, setting rule thresholds, and so on.

Note The Cisco Talos Intelligence Group (Talos) determines the appropriate state of each rule in the system-provided policies. If you use a system-provided policy as your base policy, and you allow the system to set your rules to the Secure Firewall recommended rule state, the rules in your intrusion policy match the settings recommended for your network assets.

Prerequisites for Network Analysis and Intrusion Policies

To allow the Snort inspection engine to process traffic for intrusion and malware analysis, you must have the IPS license enabled for the Firewall Threat Defense device.

You must be an Admin user to manage network analysis, intrusion policies, and perform migration tasks.

Generate New Secure Firewall Recommendations in Snort 3

Generate the Secure Firewall recommendations for the intrusion policy and then follow the steps that are listed here to create new recommended rule settings in Snort 3. Rule overheads are interpreted as **security levels** based on the threshold policies selected by you in Snort 3. Recommended action is based on the selected security level and if it is higher than the base policy, then the recommendation is not just limited to generating the events.

Prior to setting the Secure Firewall recommendations you should ask which of the three points listed below closely matches the goal:

- Increased Protection —Enable additional rules based on vulnerabilities found in the host database and do not automatically disable any rules. This will likely result in a larger rule set.
- Focused Protection—Enable additional rules and disable existing rules based on vulnerabilities found in the host database. This can increase or decrease the number of rules depending on vulnerabilities discovered.
- Higher Efficiency—Use the currently enabled rule set and disable any rules for vulnerabilities not found in the host database. This will likely result in a smaller enabled rule set.

Based on the response, the recommendation actions are as follows:

- Set recommendations to the next highest security level, and uncheck the disable rules.
- Set recommendations to the next highest security level, and check the disable rules.
- Set recommendations to the current security level, and check the disable rules.

Before you begin

Secure Firewall recommendations have the following requirements:

- Ensure that hosts are present in the system to generate recommendations.
- Protected networks configured for recommendations should map to the hosts present in the system

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Click the **Snort 3 Version** button of the intrusion policy.

Step 3 Click the **Recommendations (Not in Use)** layer to configure the rule recommendations. Click **Start**.

In the Secure Firewall Rule Recommendations window you can set the following:

- **Security Level**: Click to select the security level. Optionally, you can check the **Accept Recommendation to Disable Rules** checkbox to disable rules that are not enabled at the input security level and in protected networks. Only enable this option if you need to trim your rule set due to a high number of alerts or to improve inspection performance. The security levels are:
  - Security level 1: Connectivity Over Security

No Impact—No new rules are enabled and no existing rules are disabled. To increase the protection, select a higher security level.

Lower Security (checkbox is checked)—All rules are disabled except for the rules in the Connectivity Over Security ruleset that match potential vulnerabilities on discovered hosts. It is recommended instead to adjust the Base Policy.

- Security level 2: Balanced Security Over Connectivity

  No Impact—No new rules are enabled and no existing rules are disabled. To increase the protection, select a higher security level.

  Higher Efficiency(checkbox is checked)—Keeps existing rules that match potential vulnerabilities on discovered hosts and disables rules for vulnerabilities not found on the network.

- Security level 3: Security Over Connectivity

  Increased Security—Enables additional rules that match potential vulnerabilities on discovered hosts based on the Maximum Detection ruleset.

  Focused Security(checkbox is checked)—Enables additional rules that match vulnerabilities on discovered hosts based on the Security Over Connectivity ruleset, while disabling existing rules that do not match potential vulnerabilities on discovered hosts.

- Security level 4 : Maximum Detection

  Increased Security—Enables additional rules that match potential vulnerabilities on discovered hosts based on the Security Over Connectivity ruleset.

  Focused Security (checkbox is checked)—Enables additional rules that match vulnerabilities on discovered hosts based on the Maximum Detection ruleset, while disabling existing rules that do not match potential vulnerabilities on discovered hosts.

  Note
  Maximum Detection enables a very high number of rules and may impact performance. We recommend you to review and test this setting before deploying into a production environment.

- **Protected Networks**: Specifies the monitored networks or individual hosts to examine for recommendations. You can select one or more system or custom defined network objects from the drop-down list. By default, any IPv4 or IPv6 networks are selected, if no selection is done.

  Important
  The Secure Firewall Rule Recommendations depend on network discovery. Protected Networks apply to any hosts discovered within the ranges configured in your Network Discovery policy. For more information, see the chapter Network Discovery Policies in the *Cisco Secure Firewall Management Center Device Configuration guide*.

  Click the **Add +** button to create a new network object of type Host or Network and click **Save**.

Step 4 Generate and apply recommendations:

- **Generate**: Generates the recommendations for an intrusion policy. This action lists the rules under Recommended Rules (Not in use).
- **Generate and Apply**: Generates and applies the recommendations for an intrusion policy. This action lists the rules under Recommended Rules (In use).

Recommendations are generated successfully. A new recommendation tab appears with all the recommended rules with their corresponding recommended actions. Rule action preset filters are also available for this tab, in addition with new recommendations.

Step 5 You can verify the recommendations and then choose to apply them accordingly:

- **Accept**—Applies the previously generated recommendations for an Intrusion policy.
- **Refresh**—Regenerates and updates the rule recommendations for an Intrusion policy.
- **Edit**—It opens the Recommendations dialog box, you can provide the recommendation input values and then generate the recommendations.
- **Remove All**—Revert or remove the applied recommended rules from the policy and also removes the recommendation tab.

Under **All Rules**, there is a Recommended Rules section which displays the recommended rules.

Note
Final action for an Intrusion rule is applied based on the rule action priority order and following is the rule action priority order:

Rule Override > Generated Recommendations > Group Override > Base Policy Default Action

For enabled recommendations, Firewall Management Center considers the current state: group overrides, base policy, and recommendation configurations and priority order of actions is:

pass > block > reject > drop > rewrite > alert

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 2055.

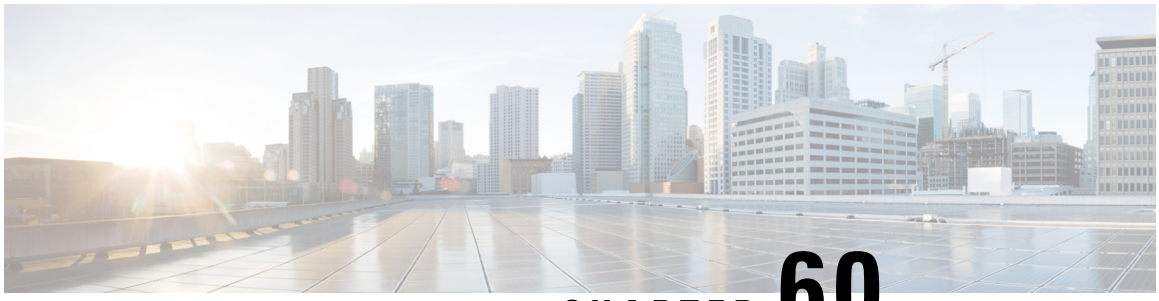

CHAPTER 60

Intrusion Prevention Performance Tuning

The following topics describe how to refine intrusion prevention performance:

- About Intrusion Prevention Performance Tuning, on page 2099
- License Requirements for Intrusion Prevention Performance Tuning, on page 2100
- Requirements and Prerequisites for Intrusion Prevention Performance Tuning, on page 2100
- Limiting Pattern Matching for Intrusions, on page 2100
- Regular Expression Limits Overrides for Intrusion Rules, on page 2101
- Overriding Regular Expression Limits for Intrusion Rules, on page 2102
- Per Packet Intrusion Event Generation Limits, on page 2103
- Limiting Intrusion Events Generated Per Packet, on page 2103
- Packet and Intrusion Rule Latency Threshold Configuration, on page 2104
- Intrusion Performance Statistic Logging Configuration, on page 2110
- Configuring Intrusion Performance Statistic Logging, on page 2110

About Intrusion Prevention Performance Tuning

Cisco provides several features for improving the performance of your system as it analyzes traffic for attempted intrusions. You can:

- specify the number of packets to allow in the event queue. You can also, before and after stream reassembly, enable or disable inspection of packets that will be rebuilt into larger streams.
- override default match and recursion limits on PCRE that are used in intrusion rules to examine packet payload content.
- elect to have the rules engine log more than one event per packet or packet stream when multiple events are generated, allowing you to collect information beyond the reported event.
- balance security with the need to maintain device latency at an acceptable level with packet and rule latency thresholding.
- configure the basic parameters of how devices monitor and report their own performance. This allows you to specify the intervals at which the system updates performance statistics on your devices.

You configure these performance settings on a per-access-control-policy basis, and they apply to all intrusion policies invoked by that parent access control policy.

Note Snort 2 is not supported on threat defense Version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the Firewall Management Center guide that matches your Firewall Threat Defense version.

License Requirements for Intrusion Prevention Performance Tuning

Threat Defense License

IPS

Requirements and Prerequisites for Intrusion Prevention Performance Tuning

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Limiting Pattern Matching for Intrusions

Procedure

Step 1 In the access control policy editor, click **Advanced** (**Policies** > **Access Control heading** > **Access Control**, click **Edit** and then click **Advanced Settings**).

In the new UI, select **Advanced Settings** from the drop-down arrow at the end of the packet flow line.

Step 2 Click **Edit** (✎) next to **Performance Settings**.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Click **Pattern Matching Limits** in the **Performance Settings** pop-up window.

Step 4 Enter a value for the maximum number of events to queue in the **Maximum Pattern States to Analyze Per Packet** field.

Step 5 To disable the inspection of packets that will be rebuilt into larger streams of data before and after stream reassembly in Snort 2, check the **Disable Content Checks on Traffic Subject to Future Reassembly** check box. Inspection before and after reassembly requires more processing overhead and may decrease performance.

Important
In Snort 3, the **Disable Content Checks on Traffic Subject to Future Reassembly** check box settings are:

- Checked—Indicates detecting TCP payload before reassembly. It includes inspection of packets before and after stream reassembly. This process requires more processing overhead and may decrease performance.
- Unchecked—Indicates detecting TCP payload after reassembly.

Step 6 Click **OK**.

Step 7 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Regular Expression Limits Overrides for Intrusion Rules

The default regular expression limits ensure a minimum level of performance. Overriding these limits could increase security, but could also significantly impact performance by permitting packet evaluation against inefficient regular expressions.

Caution Do not override default PCRE limits unless you are an experienced intrusion rule writer with knowledge of the impact of degenerative patterns.

Table 113: Regular Expression Constraint Options

Option	Description
Match Limit State	Specifies whether to override **Match Limit**. You have the following options: • select **Default** to use the value configured for **Match Limit** • select **Unlimited** to permit an unlimited number of attempts • select **Custom** to specify either a limit of 1 or greater for **Match Limit**, or to specify 0 to completely disable PCRE match evaluations

Option	Description
Match Limit	Specifies the number of times to attempt to match a pattern defined in a PCRE regular expression.
Match Recursion Limit State	Specifies whether to override **Match Recursion Limit**. You have the following options: • select **Default** to use the value configured for **Match Recursion Limit** • select **Unlimited** to permit an unlimited number of recursions • select **Custom** to specify either a limit of 1 or greater for **Match Recursion Limit**, or to specify 0 to completely disable PCRE recursions Note that for **Match Recursion Limit** to be meaningful, it must be smaller than **Match Limit**.
Match Recursion Limit	Specifies the number of recursions when evaluating a PCRE regular expression against the packet payload.

Related Topics

Overview: The pcre Keyword

Overriding Regular Expression Limits for Intrusion Rules

Procedure

Step 1 In the access control policy editor, click **Advanced**.

In the new UI, select **Advanced Settings** from the drop-down arrow at the end of the packet flow line.

Step 2 Click **Edit** (✎) next to **Performance Settings**.

If **View** (👁) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Click **Regular Expression Limits** in the **Performance Settings** pop-up window.

Step 4 You can modify any of the options as described in Regular Expression Limits Overrides for Intrusion Rules, on page 2101.

Step 5 Click **OK**.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Per Packet Intrusion Event Generation Limits

When the intrusion rules engine evaluates traffic against rules, it places the events generated for a given packet or packet stream in an event queue, then reports the top events in the queue to the user interface. When configuring the intrusion event logging limits, you can specify how many events can be placed in the queue and how many are logged, and select the criteria for determining event order within the queue.

Table 114: Intrusion Event Logging Limits Options

Option	Description
Maximum Events Stored Per Packet	The maximum number of events that can be stored for a given packet or packet stream.
Maximum Events Logged Per Packet	The number of events logged for a given packet or packet stream. This cannot exceed the **Maximum Events Stored Per Packet** value.
Prioritize Event Logging By	The value used to determine event ordering within the event queue. The highest ordered event is reported through the user interface. You can select from: • `priority`, which orders events in the queue by the event priority. • `content_length`, which orders events by the longest identified content match. When events are ordered by content length, rule events always take precedence over decoder and preprocessor events.

Limiting Intrusion Events Generated Per Packet

Procedure

Step 1 In the access control policy editor, click **Advanced**.

In the new UI, select **Advanced Settings** from the drop-down arrow at the end of the packet flow line.

Step 2 Click **Edit** (✎) next to **Performance Settings**.

If **View** (👁) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Click **Intrusion Event Logging Limits** in the **Performance Settings** pop-up window.

Step 4 You can modify any of the options in Per Packet Intrusion Event Generation Limits, on page 2103.

Step 5 Click **OK**.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Packet and Intrusion Rule Latency Threshold Configuration

Each access control policy has latency-based settings that use thresholding to manage packet and rule processing performance.

Packet latency thresholding measures the total elapsed time taken to process a packet by applicable decoders, preprocessors, and rules, and ceases inspection of the packet if the processing time exceeds a configurable threshold.

Rule latency thresholding measures the elapsed time each rule takes to process an individual packet, suspends the violating rule along with a group of related rules for a specified time if the processing time exceeds the rule latency threshold a configurable consecutive number of times, and restores the rules when the suspension expires.

Note Snort 2 is not supported on threat defense Version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the Firewall Management Center guide that matches your Firewall Threat Defense version.

Latency-Based Performance Settings

By default, the system takes latency-based performance settings from the latest intrusion rule update deployed on your system.

The latency settings that are actually applied depend on the security level of the network analysis policy (NAP) associated with the access control policy. Generally, this is the default NAP policy. However, if custom network analysis rules are configured, and if any of these specify a NAP policy that is more secure than the default NAP policy, then latency settings are based on the most secure NAP policy among the custom rules. If the default NAP policy or any custom rules invoke a custom NAP policy, then the security level used in the evaluation is the system-provided base policy on which each custom NAP policy is based.

The above is true regardless of whether the effective threshold and/or network analysis configurations are inherited or configured directly in the policy.

Packet Latency Thresholding

Packet latency thresholding measures elapsed time, not just processing time, in order to more accurately reflect the actual time required for the rule to process a packet. However, latency thresholding is a software-based latency implementation that does not enforce strict timing.

The trade-off for the performance and latency benefits derived from latency thresholding is that uninspected packets could contain attacks. A timer starts for each packet when decoder processing begins. Timing continues either until all processing ends for the packet or until the processing time exceeds the threshold at a timing test point.

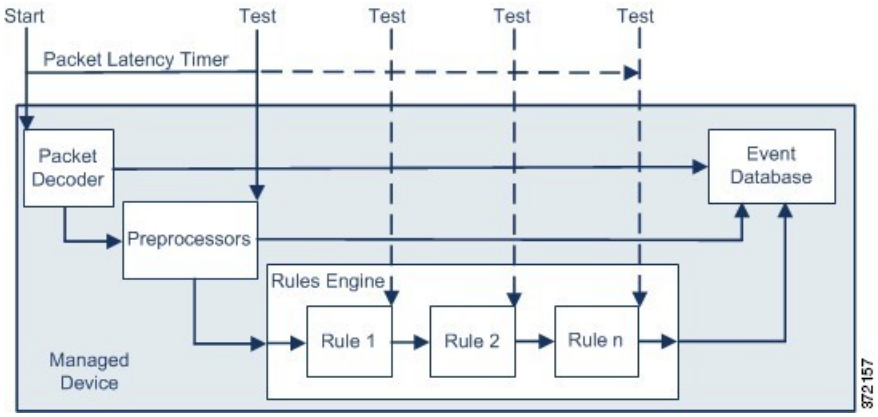

As illustrated in the above figure, packet latency timing is tested at the following test points:

- after the completion of all decoder and preprocessor processing and before rule processing begins
- after processing by each rule

If the processing time exceeds the threshold at any test point, packet inspection ceases.

Tip Total packet processing time does not include routine TCP stream or IP fragment reassembly times.

Packet latency thresholding has no effect on events triggered by a decoder, preprocessor, or rule processing the packet. Any applicable decoder, preprocessor, or rule triggers normally until a packet is fully processed, or until packet processing ends because the latency threshold is exceeded, whichever comes first. If a drop rule detects an intrusion in an inline deployment, the drop rule triggers an event and the packet is dropped.

Note No packets are evaluated against rules after processing for that packet ceases because of a packet latency threshold violation. A rule that would have triggered an event cannot trigger that event, and for drop rules, cannot drop the packet.

Packet latency thresholding can improve system performance in both passive and inline deployments, and can reduce latency in inline deployments, by stopping inspection of packets that require excessive processing time. These performance benefits might occur when, for example:

- for both passive and inline deployments, sequential inspection of a packet by multiple rules requires an excessive amount of time
- for inline deployments, a period of poor network performance, such as when someone downloads an extremely large file, slows packet processing

In a passive deployment, stopping the processing of packets might not contribute to restoring network performance because processing simply moves to the next packet.

Packet Latency Thresholding Notes

By default, the latency-based performance settings for packet handling is disabled. You may choose to enable it. However, Cisco recommends that you do not change the default value for the threshold setting.

The information in this below applies only if you choose to specify custom values.

Table 115: Packet Latency Thresholding Option

Option	Description
Threshold (microseconds)	Specifies the time, in microseconds, when inspection of a packet ceases.

Enabling Packet Latency Thresholding

Procedure

Step 1 In the access control policy editor, click **Advanced**.

In the new UI, select **Advanced Settings** from the drop-down arrow at the end of the packet flow line.

Step 2 Click **Edit** () next to **Latency-Based Performance Settings**.

If **View** () appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings.

Step 3 Click **Packet Handling** in the **Latency-Based Performance Settings** pop-up window.

Step 4 Check the **Enabled** check box.

Step 5 Click **OK**.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Configuring Packet Latency Thresholding

By default, the latency-based performance settings for packet handling is disabled. You may choose to enable it. However, Cisco recommends that you do not change the default value for the threshold setting.

Procedure

Step 1 In the access control policy editor, click **Advanced**.

In the new UI, select **Advanced Settings** from the drop-down arrow at the end of the packet flow line.

Step 2 Click **Edit** () next to **Latency-Based Performance Settings**.

System () > **Monitoring** > **Statistics**

Step 3 If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 4 Click **Packet Handling** in the **Latency-Based Performance Settings** pop-up window.

By default, **Installed Rule Update** is selected. We recommend using this default.

The values displayed do not reflect the automated settings.

Step 5 If you choose to specify custom values:

- Check the **Enabled** check box, and see Packet Latency Thresholding Notes, on page 2105 for recommended minimum **Threshold** settings.
- You must specify custom values in both the packet handling tab and the rule handling tab.

Step 6 Click **OK**.

Step 7 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Rule Latency Thresholding

Rule latency thresholding measures elapsed time, not just processing time, in order to more accurately reflect the actual time required for the rule to process a packet. However, latency thresholding is a software-based latency implementation that does not enforce strict timing.

The trade-off for the performance and latency benefits derived from latency thresholding is that uninspected packets could contain attacks. A timer measures the processing time each time a packet is processed against a group of rules. Any time the rule processing time exceeds a specified rule latency threshold, the system increments a counter. If the number of consecutive threshold violations reaches a specified number, the system takes the following actions:

- suspends the rules for the specified period
- triggers an event indicating the rules have been suspended
- re-enables the rules when the suspension expires
- triggers an event indicating the rules have been re-enabled

The system zeroes the counter when the group of rules has been suspended, or when rule violations are not consecutive. Permitting some consecutive violations before suspending rules lets you ignore occasional rule violations that might have negligible impact on performance and focus instead on the more significant impact of rules that repeatedly exceed the rule latency threshold.

The following example shows five consecutive rule processing times that do not result in rule suspension.

In the above example, the time required to process each of the first three packets violates the rule latency threshold of 1000 microseconds, and the violations counter increments with each violation. Processing of the fourth packet does not violate the threshold, and the violations counter resets to zero. The fifth packet violates the threshold and the violations counter restarts at one.

The following example shows five consecutive rule processing times that do result in rule suspension.

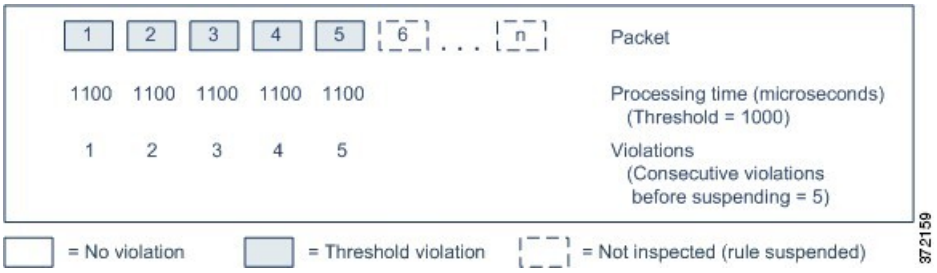

In the second example, the time required to process each of the five packets violates the rule latency threshold of 1000 microseconds. The group of rules is suspended because the rule processing time of 1100 microseconds for each packet violates the threshold of 1000 microseconds for the specified five consecutive violations. Any subsequent packets, represented in the figure as packets 6 through n, are not examined against suspended rules until the suspension expires. If more packets occur after the rules are re-enabled, the violations counter begins again at zero.

Rule latency thresholding has no effect on intrusion events triggered by the rules processing the packet. A rule triggers an event for any intrusion detected in the packet, regardless of whether the rule processing time exceeds the threshold. If the rule detecting the intrusion is a drop rule in an inline deployment, the packet is dropped. When a drop rule detects an intrusion in a packet that results in the rule being suspended, the drop rule triggers an intrusion event, the packet is dropped, and that rule and all related rules are suspended.

Note Packets are not evaluated against suspended rules. A suspended rule that would have triggered an event cannot trigger that event and, for drop rules, cannot drop the packet.

Rule latency thresholding can improve system performance in both passive and inline deployments, and can reduce latency in inline deployments, by suspending rules that take the most time to process packets. Packets are not evaluated again against suspended rules until a configurable time expires, giving the overloaded device time to recover. These performance benefits might occur when, for example:

- hastily written, largely untested rules require an excessive amount of processing time
- a period of poor network performance, such as when someone downloads an extremely large file, causes slow packet inspection

Rule Latency Thresholding Notes

By default, latency-based performance settings for both packet and rule handling are automatically populated by the latest deployed intrusion rule update, and we recommend that you do not change the default.

The information in this topic applies only if you choose to specify custom values.

Rule latency thresholding suspends rules for the time specified by **Suspension Time** when the time rules take to process a packet exceeds **Threshold** for the consecutive number of times specified by **Consecutive Threshold Violations Before Suspending Rule**.

You can enable rule 134:1 to generate an event when rules are suspended, and rule 134:2 to generate an event when suspended rules are enabled.

Table 116: Rule Latency Thresholding Options

Option	Description
Threshold	Specifies the time in microseconds that rules should not exceed when examining a packet.
Consecutive Threshold Violations Before Suspending Rule	Specifies the consecutive number of times rules can take longer than the time set for **Threshold** to inspect packets before rules are suspended.
Suspension Time	Specifies the number of seconds to suspend a group of rules.

Configuring Rule Latency Thresholding

By default, latency-based performance settings for both packet and rule handling are automatically populated by the latest deployed intrusion rule update, and we recommend that you do not change the default.

Procedure

Step 1 In the access control policy editor, click **Advanced**.

In the new UI, select **Advanced Settings** from the drop-down arrow at the end of the packet flow line.

Step 2 Click **Edit** (✎) next to **Latency-Based Performance Settings**.

If **View** (👁) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Click **Rule Handling** in the **Latency-Based Performance Settings** pop-up window.

By default, **Installed Rule Update** is selected. We recommend using this default.

The values displayed do not reflect the automated settings.

Step 4 If you choose to specify custom values:

- You can configure any of the options in Rule Latency Thresholding Notes, on page 2108.
- You must specify custom values in both the packet handling tab and the rule handling tab.

Step 5 Click **OK**.

Step 6 Click **Save** to save the policy.

What to do next

- If you want to generate events, enable latency rules 134:1 and 134:2.
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Intrusion Performance Statistic Logging Configuration

Sample time (seconds) and Minimum number of packets

When the number of seconds specified elapses between performance statistics updates, the system verifies it has analyzed the specified number of packets. If it has, the system updates performance statistics. Otherwise, the system waits until it analyzes the specified number of packets.

Caution Configuring a very low value (for example 1 second) for the sample time can cause a huge impact on the device; the performance statistics logged on the device can cause disk space issues and affect the operation of the device. Hence we recommend you do not configure a very low value.

Troubleshooting Options: Log Session/Protocol Distribution

Support might ask you during a troubleshooting call to log protocol distribution, packet length, and port statistics.

Caution Do not enable **Log Session/Protocol Distribution** unless instructed to by Support.

Troubleshooting Options: Summary

Support might ask you during a troubleshooting call to configure the system to calculate the performance statistics only when the Snort process is shut down or restarted. To enable this option, you must also enable the **Log Session/Protocol Distribution** troubleshooting option.

Caution Do not enable **Summary** unless instructed to do so by Support.

Configuring Intrusion Performance Statistic Logging

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** (✎) next to **Performance Settings**.

In the new UI, select **Advanced Settings** from the drop-down arrow at the end of the packet flow line.

If **View** (👁) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 Click **Performance Statistics** in the pop-up window that appears.

Step 3 Modify the **Sample time** or **Minimum number of packets** as described in Intrusion Performance Statistic Logging Configuration, on page 2110.

Caution
Configuring a very low value (for example 1 second) for the **Sample time** can cause a huge impact on the device; the performance statistics logged on the device can cause disk space issues and affect the operation of the device. Hence we recommend you do not configure a very low value.

Step 4 Optionally, expand the **Troubleshoot Options** section and modify those options only if asked to do so by Support.

Step 5 Click **OK**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

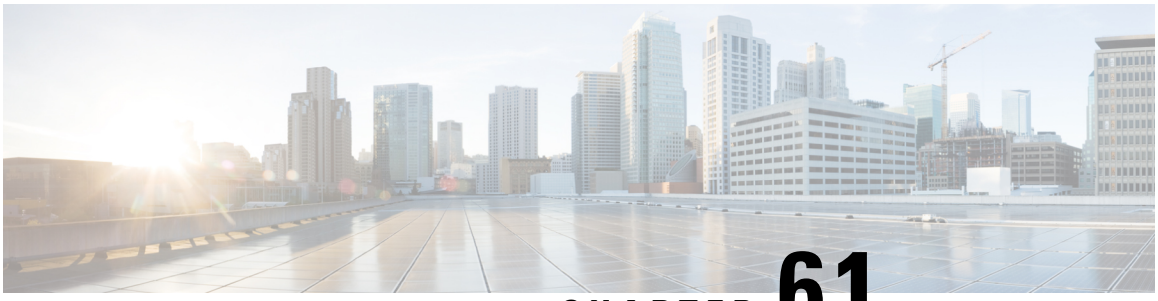

CHAPTER 61

Get Started with Snort 3 Network Analysis Policies

This chapter provides an insight into network analysis policy basics, prerequisites, and how to manage network analysis policies. It also provides information on custom network analysis policy creation and network analysis policy settings.

- Overview of Network Analysis Policies, on page 2113
- Manage Network Analysis Policies, on page 2114
- Snort 3 Definitions and Terminologies for Network Analysis Policy, on page 2115
- Prerequisites for Network Analysis and Intrusion Policies, on page 2117
- Custom Network Analysis Policy Creation for Snort 3, on page 2117
- Network Analysis Policy Settings and Cached Changes, on page 2145

Overview of Network Analysis Policies

Network analysis policies govern many traffic preprocessing options, and are invoked by advanced settings in your access control policy. Network analysis-related preprocessing occurs after Security Intelligence matching and SSL decryption, but before intrusion or file inspection begins.

By default, the system uses the *Balanced Security and Connectivity* network analysis policy to preprocess all traffic handled by an access control policy. However, you can choose a different default network analysis policy to perform this preprocessing. For your convenience, the system provides a choice of several non-modifiable network analysis policies, which are tuned for a specific balance of security and connectivity by the Cisco Talos Intelligence Group (Talos). You can also create a custom network analysis policy with custom preprocessing settings.

Tip System-provided intrusion and network analysis policies are similarly named but contain different configurations. For example, the Balanced Security and Connectivity network analysis policy and the Balanced Security and Connectivity intrusion policy work together and can both be updated in intrusion rule updates. However, the network analysis policy governs mostly preprocessing options, whereas the intrusion policy governs mostly intrusion rules. Network analysis and intrusion policies work together to examine your traffic.

You can also tailor traffic preprocessing options to specific security zones, networks, and VLANs by creating multiple custom network analysis policies, then assigning them to preprocess different traffic. (Note that ASA FirePOWER cannot restrict preprocessing by VLAN.)

Manage Network Analysis Policies

Under your user name in the toolbar, the system displays a tree of available domains. To switch domains, choose the domain you want to access.

Note Snort 2 is not supported on threat defense Version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the management center guide that matches your threat defense version.

Procedure

Step 1 Choose one of the following paths to access the network analysis policy.

- **Policies** > **Access Control heading** > **Access Control**, and then click **Network Analysis Policy**
- **Policies** > **Access Control heading** > **Intrusion**, and then click **Network Analysis Policies**

Note
If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Manage your network analysis policy:

- Create—If you want to create a new network analysis policy, click **Create Policy**.

  Two versions of the network analysis policy are created, a **Snort 2 Version** and a **Snort 3 Version**.

  - For the Snort 2 version, see *Custom Network Analysis Policy Creation for Snort 2* in the *Cisco Secure Firewall Management Center Configuration Guide*.
  - For the Snort 3 version, see Custom Network Analysis Policy Creation for Snort 3, on page 2117.

- Delete—If you want to delete a network analysis policy, click the **Delete** icon, then confirm that you want to delete the policy. You cannot delete a network analysis policy if an access control policy references it.

  If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

- Edit—If you want to edit an existing network analysis policy, click the **Edit** icon.

  If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

- Report—Click the **Report** icon; see *Generating Current Policy Reports* in the *Cisco Secure Firewall Management Center Configuration Guide*.

Snort 3 Definitions and Terminologies for Network Analysis Policy

The following table lists the Snort 3 concepts and terms used in the Network Analysis Policy.

Table 117: Snort 3 Definitions and Terminologies for Network Analysis Policy

Term	Description
Inspectors	Inspectors are plugins that process packets (similar to the Snort 2 preprocessor).
Binder inspector	Binder inspector defines the flow when a particular inspector has to be accessed and taken into consideration. When the traffic matches the conditions defined in the binder inspector, only then do the values/configurations for that inspector come into effect. For more information, see *Binder Inspector* in Custom Network Analysis Policy Creation for Snort 3, on page 2117.
Singleton inspectors	Singleton inspectors contain one instance. These inspectors do not support adding more instances like multiton inspectors. Settings of singleton inspector are applied to the entire traffic matching that inspector and not to a specific traffic segment. For more information, see *Singleton Inspectors* in Custom Network Analysis Policy Creation for Snort 3, on page 2117.
Multiton inspectors	Multiton inspectors contain multiple instances which you can configure as needed. These inspectors support configuring settings based on specific conditions, such as network, port, and VLAN. One set of supported settings is called an instance. For more information, see *Multiton Inspectors* in Custom Network Analysis Policy Creation for Snort 3, on page 2117.

Term	Description
Schema	The schema file is based on the OpenAPI JSON specification, and it validates the content that you upload or download. You can download the schema file and open it using any third-party JSON editor, such as Swagger editor. The schema file helps you to identify what parameters can be configured for inspectors with their corresponding allowed values, range, and accepted patterns to be used. For more information, see Customize the Network Analysis Policy, on page 2125.
Sample file	It is a pre-existing template that contains example configurations to help you with configuring the inspectors. You can refer to the example configurations included in the sample file and make any changes that you may require. For more information, see Customize the Network Analysis Policy, on page 2125.
Full configuration	You can download the entire inspector configurations in a single file. All information regarding the inspector configuration is available in this file. The full configuration is a merged configuration of the default configuration (rolled out as a part of the LSP updates by Cisco Talos) and the custom NAP inspector configurations. For more information, see Customize the Network Analysis Policy, on page 2125.

Term	Description
Overridden configuration	In the **Snort 3 Version** of the network analysis policy page: • Under **Actions** > **Upload**, you can click **Overridden Configuration** to upload the JSON file that contains the overridden configuration. • Under **Actions** > **Download**, you can click **Overridden Configuration** to download the inspector configuration that has been overridden. If you have not overridden any inspector configuration, then this option is disabled. When you override the inspector configuration, then this option is enabled automatically to allow you to download. For more information, see Customize the Network Analysis Policy, on page 2125.

Related Topics

Custom Network Analysis Policy Creation for Snort 3, on page 2117
Customize the Network Analysis Policy, on page 2125
Network Analysis Policy Mapping, on page 2122

Prerequisites for Network Analysis and Intrusion Policies

To allow the Snort inspection engine to process traffic for intrusion and malware analysis, you must have the IPS license enabled for the Firewall Threat Defense device.

You must be an Admin user to manage network analysis, intrusion policies, and perform migration tasks.

Custom Network Analysis Policy Creation for Snort 3

The default network analysis policy is tuned for typical network requirements and optimal performance. Usually, the default network analysis policy suffices most network requirements and you might not need to customize the policy. However, when you have a specific network requirement or when you are facing performance issues, the default network analysis policy can be customized. Note that customizing the network analysis policy is an advanced configuration that should be done only by advanced users or Cisco support.

Network analysis policy configuration for Snort 3 is a data-driven model, which is based on JSON and JSON Schema. Schema is based on the OpenAPI specification, and it helps you get a view of the supported inspectors, settings, settings type, and valid values. The Snort 3 inspectors are plugins that process packets (similar to the Snort 2 preprocessor). Network analysis policy configuration is available to download in the JSON format.

In Snort 3, the list of inspectors and settings are not in a one-to-one mapping with the Snort 2 list of preprocessors and settings. Also, the number of inspectors and settings available in Firewall Management Center is a subset of the inspectors and settings that Snort 3 supports. See https://snort.org/snort3 for more

information on Snort 3. See https://www.cisco.com/go/snort3-inspectors for more information on the inspectors available in Firewall Management Center.

Note

- While upgrading the Firewall Management Center to the 7.0 release, the changes that were done in the Snort 2 version of the network analysis policy are not migrated to Snort 3 after the upgrade.
- Unlike the intrusion policy, there is no option to synchronize Snort 2 network analysis policy settings to Snort 3.

Default Inspector Updates

Lightweight Security Package (LSP) updates may contain new inspectors or modifications to integer ranges for existing inspector configurations. Following the installation of an LSP, new inspectors and/or updated ranges will be available under **Inspectors** in the **Snort 3 Version** of your network analysis policy.

Binder Inspector

Binder inspector defines the flow when a particular inspector has to be accessed and taken into consideration. When the traffic matches the conditions defined in the binder inspector, only then the values/configurations for that inspector come into effect. For example:

For the *imap* inspector, the binder defines the following condition when it has to be accessed. That is when:

- Service is equal to imap.
- Role is equal to any.

If these conditions are met, then use the type imap.

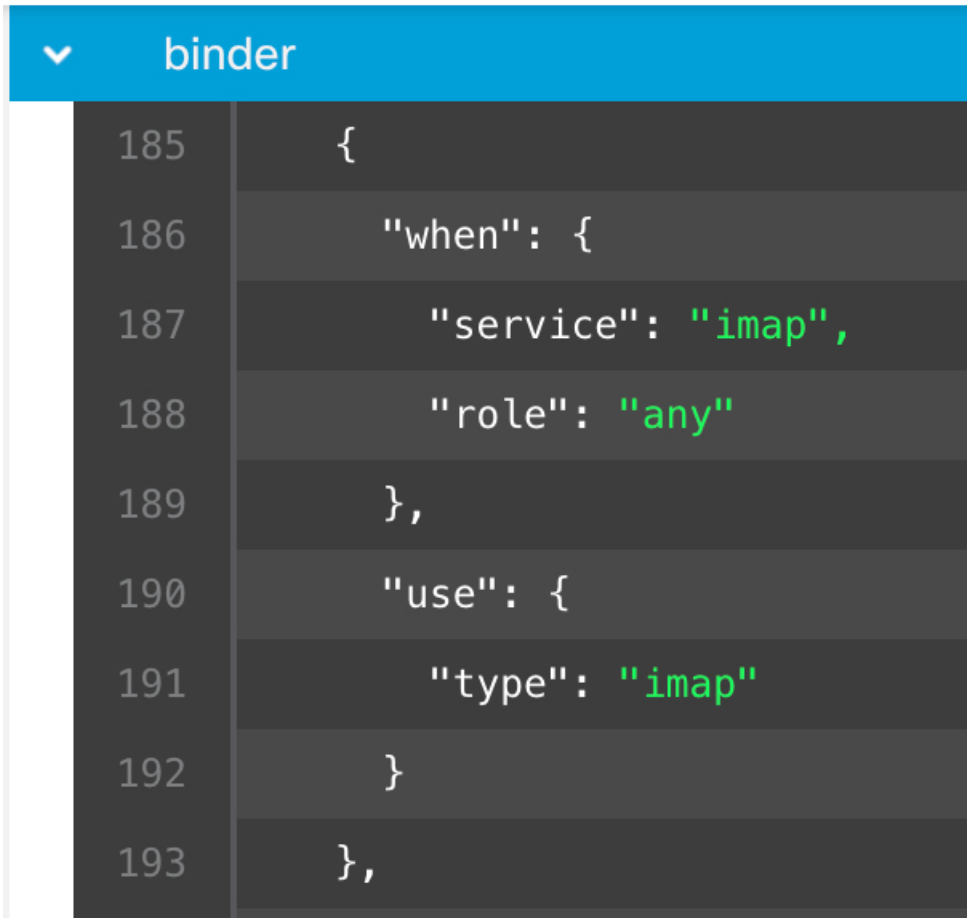

Singleton Inspectors

Singleton inspectors contain a single instance. These inspectors do not support adding more instances like multiton inspectors. Settings of singleton inspector are applied to the entire traffic and not to a specific traffic segment.

For example:

```
{
   "normalizer":{
      "enabled":true,
      "type":"singleton",
      "data":{
         "ip4":{
            "df":true
         }
      }
   }
}
```

Multiton Inspectors

Multiton inspectors contain multiple instances which you can configure as needed. These inspectors support configuring settings based on specific conditions, such as network, port, and VLAN. One set of supported

settings is called an instance. There is a default instance, and you can also add additional instances based on specific conditions. If the traffic matches that condition, the settings from that instance are applied. Otherwise, the settings from the default instance are applied. Also, the name of the default instance is the same as the inspector's name.

For a multiton inspector, when you upload the overridden inspector configuration, you also need to include/define a matching binder condition (conditions under when the inspector has to be accessed or used) for each instance in the JSON file, otherwise, the upload will result in an error. You can also create new instances, but make sure that you include a binder condition for every new instance that you create to avoid errors.

For example:

- Multiton inspector where the default instance is modified.

```
{
   "http_inspect":{
      "enabled":true,
      "type":"multiton",
      "instances":[
         {
            "name":"http_inspect",
            "data":{
               "response_depth":5000
            }
         }
      ]
   }
}
```

- Multiton inspector where the default instance and default binder is modified.

```
{
   "http_inspect":{
      "enabled":true,
      "type":"multiton",
      "instances":[
         {
            "name":"http_inspect",
            "data":{
               "response_depth":5000
            }
         }
      ]
   },
   "binder":{
      "type":"binder",
      "enabled":true,
      "rules":[
         {
            "use":{
               "type":"http_inspect"
            },
            "when":{
               "role":"any",
               "ports":"8080",
               "proto":"tcp",
               "service":"http"
            }
         }
      ]
   }
}
```

- Multiton inspector where a custom instance and a custom binder is added.

```
{
   "http_inspect":{
      "enabled":true,
      "type":"multiton",
      "instances":[
         {
            "name":"http_inspect1",
            "data":{
               "response_depth":5000
            }
         }
      ]
   },
   "binder":{
      "type":"binder",
      "enabled":true,
      "rules":[
         {
            "use":{
               "type":"http_inspect",
               "name":"http_inspect1"
            },
            "when":{
               "role":"any",
               "ports":"8080",
               "proto":"tcp",
               "service":"http"
            }
         }
      ]
   }
}
```

Common Industrial Protocol Safety

Common Industrial Protocol (CIP) Safety is a set of extensions to the CIP that enables the safe operation of devices. It also provides fail-safe communication between different nodes on a CIP network.

The CIP Safety protocol comprises two main components:

- CIP Safety segments—Used in Forward Open messages to exchange safety parameters for the subsequent safety session.
- CIP Safety messages—Used to exchange actual safety information.

The CIP inspector detects and identifies:

- CIP as a service and client
- Payloads, such as CIP Read, CIP Admin, CIP Infrastructure, and CIP Write

  The CIP inspector can parse the CIP segments and detect the CIP Safety segments in the Forward Open requests.

To test the CIP Safety feature, you must enable the CIP inspector. See Detect and Block Safety Segments in CIP Packets, on page 2122.

Detect and Block Safety Segments in CIP Packets

Use case: To detect and block CIP safety segments while allowing other CIP packets:

- Create a custom network analysis policy called **cip_safety**.
- Create access control rules in your access control policy to block CIP Safety and to allow all other packets.

To test the CIP Safety feature, enable the CIP inspector in the management center and assign it to an access control policy.

Procedure

Step 1 Go to **Policies** > **Intrusion** > **Network Analysis Policies**.

Step 2 Click the **Snort 3 Version** of the network analysis policy **cip_safety** that you created.

Step 3 Under **Inspectors**, click **cip** to expand it.

The default configuration appears in the left column and the overridden configuration appears in the right column under the inspector.

Step 4 Under **Overridden Configuration** on the right column, click the **Edit Inspector** icon and change the "enabled" field in **cip** from false (default) to true.

Step 5 Click **OK**.

Step 6 Click **Save**.

Step 7 To assign the **cip** inspector to the access control policy, choose **Policies** > **Access Control** > **Edit** and choose the **Advanced Settings** option from the **More** drop-down arrow at the end of the packet flow line.

Step 8 Click **Edit** (✎) next to **Network Analysis and Intrusion Policies**.

Step 9 In the **Network Analysis and Intrusion Policies** window, choose the access control policy **cip_safety** that you created from the **Default Network Analysis Policy** drop-down list.

The CIP inspector is now enabled in the management center and you can create the custom access control rules to block CIP Safety and to allow all other CIP packets.

Step 10 After you send live traffic containing CIP Safety packet flows, go to **Connection Events** to verify that the payload is the expected payload that contains CIP Safety packet logs for the detection and block use case as mentioned in this procedure. **CIP** is detected as an application protocol and client (see the **Application Protocol** and **Client** fields), and **CIP Safety** is shown under the **Web Application** field.

Network Analysis Policy Mapping

For network analysis policies, Cisco Talos provides mapping information, which is used to find the corresponding Snort 2 version of the policies for the Snort 3 version.

This mapping ensures that the Snort 3 version of policies has its equivalent Snort 2 version.

View Network Analysis Policy Mapping

Procedure

Step 1 Go to **Policies** > **Intrusion** > **Network Analysis Policies**.

Step 2 Click **NAP Mapping**.

Step 3 Expand the arrow for **View Mappings**.

The Snort 3 network analysis policies that are automatically mapped to a Snort 2 equivalent policy are displayed.

Step 4 Click **OK**.

Create a Network Analysis Policy

All the existing network analysis policies are available in Firewall Management Center with their corresponding Snort 2 and Snort 3 versions. When you create a new network analysis policy, it is created with both the Snort 2 version and the Snort 3 version.

Procedure

Step 1 Go to **Policies** > **Intrusion** > **Network Analysis Policies**.

Step 2 Click **Create Policy**.

Step 3 Enter the **Name** and **Description**.

Step 4 Select a **Base Policy** and click **Save**.

The new network analysis policy is created with its corresponding **Snort 2 Version** and **Snort 3 Version**.

Modify the Network Analysis Policy

You can modify the network analysis policy to change its name, description, or the base policy.

Procedure

Step 1 Go to **Policies** > **Intrusion** > **Network Analysis Policies**.

Step 2 Click **Edit** to change the name, description, inspection mode, or the base policy.

Attention
Detection mode deprecation: From management center 7.4.0 onwards, for a network analysis policy (NAP), the **Detection** inspection mode is deprecated and will be removed in an upcoming release.

The **Detection** mode was intended to be used as a test mode so that you can enable inspections and see how they behave in your network before setting it to drop traffic, that is, to show traffic that would be dropped.

This behavior is improved where all inspector drops are controlled by the rule state, and you can set each one to generate events. This is done to test them before configuring the rule state to drop traffic. As we now have granular control over traffic drops in Snort 3, the **Detection** mode only adds more complexity to the product and is not needed, so the detection mode is deprecated.

If you change a NAP in **Detection** mode to **Prevention**, the NAP that processes the traffic of intrusion events and have the result "will be dropped" will now be "dropped" and the corresponding traffic will drop the traffic from these events. This is applicable for rules whose GIDs are not 1 or 3. GIDs 1 and 3 are text/compiled rules (typically provided by Talos or from your custom/imported rules) and all other GIDs are inspections for anomalies. These are more uncommon rules to trigger in a network. Changing to **Prevention** mode is unlikely to have any impact on the traffic. You need to just disable the intrusion rule that is applicable for the dropped traffic and set it to just generate or disable.

We recommend you choose **Prevention** as the inspection mode, but if you choose **Prevention**, you cannot revert to **Detection** mode.

Note
If you edit the network analysis policy name, description, base policy, and inspection mode, the edits are applied to both the Snort 2 and Snort 3 versions. If you want to change the inspection mode for a specific version, then you can do that from within the network analysis policy page for that respective version.

Step 3 Click **Save**.

Search for an Inspector on the Network Analysis Policy Page

On the Snort 3 version of the network analysis policy page, you may need to search for an inspector by entering any relevant text in the search bar.

Procedure

Step 1 Go to **Policies** > **Intrusion** > **Network Analysis Policies**.

Step 2 Go to the **Snort 3 Version** of the network analysis policy.

Step 3 Enter an inspector's name or any relevant text to search for in the **Search** bar.

All the inspectors matching the text you search for are displayed.

For example, if you enter `pop`, then the pop inspector and the binder inspector are shown as matching results on the screen.

Related Topics

Examples of Custom Network Analysis Policy Configuration, on page 2134
View the List of Inspectors with Overrides, on page 2130
Snort 3 Definitions and Terminologies for Network Analysis Policy, on page 2115
Customize the Network Analysis Policy, on page 2125
Make Inline Edit for an Inspector to Override Configuration, on page 2129

Copy the Inspector Configuration

You can copy the inspector configuration for the Snort 3 version of the network analysis policy according to your requirements.

Procedure

Step 1 Go to **Policies** > **Intrusion** > **Network Analysis Policies**.

Step 2 Go to the **Snort 3 Version** of the network analysis policy.

Step 3 Under **Inspectors**, expand the required inspector for which you want to copy the configuration.

The default configuration is displayed on the left column and the overridden configuration is displayed on the right column under the inspector.

Step 4 Click the **Copy to clipboard** icon to copy the inspector configuration to the clipboard for one or both of the following.

- **Default Configuration** in the left column
- **Overridden Configuration** in the right column

Step 5 Paste the copied inspector configuration to a JSON editor to make any edits you may require.

Related Topics

Customize the Network Analysis Policy, on page 2125

Customize the Network Analysis Policy

You can customize the Snort 3 version of the network analysis policy according to your requirements.

Procedure

Step 1 Go to **Policies** > **Intrusion** > **Network Analysis Policies**.

Step 2 Go to the **Snort 3 Version** of the network analysis policy.

Step 3 Click the **Actions** drop-down menu.

The following options are displayed:

- View Schema
- Download Schema / Download Sample File / Template
- Download Full Configuration
- Download Overridden Configuration
- Upload Overridden Configuration

Step 4 Click **View Schema** to open the schema file directly in a browser.

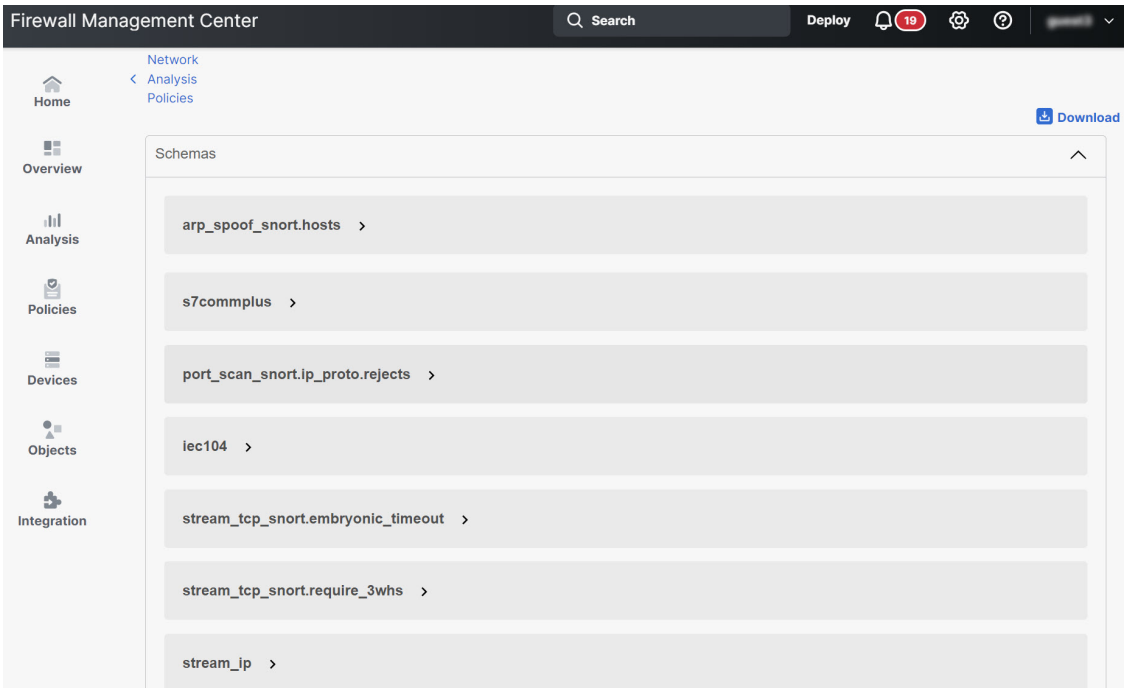

Step 5 You can download the schema file, sample file / template, full configuration, or overridden configuration as needed.

These options provide you an insight about the allowed values, range, and patterns, existing and default inspector configurations, and overridden inspector configurations.

a) Click **Download Schema** to download the schema file.

The schema file validates the content that you upload or download. You can download the schema file and open it using any third-party JSON editor. The schema file helps you to identify what parameters can be configured for inspectors with their corresponding allowed values, range, and accepted patterns to be used.

For example, for the *arp_spoof_snort* inspector, you can configure the hosts. The hosts include the *mac* and *ip* address values. The schema file shows the following accepted pattern for these values.

- **mac** – `pattern: ^([0-9A-Fa-f]{2}[:-]){5}([0-9A-Fa-f]{2})$`
- **ip** – `pattern: ^([0-9]{1,3}.[0-9]{1,3}.[0-9]{1,3}.[0-9]{1,3}(/[0-9]{1,2}){0,1})$`

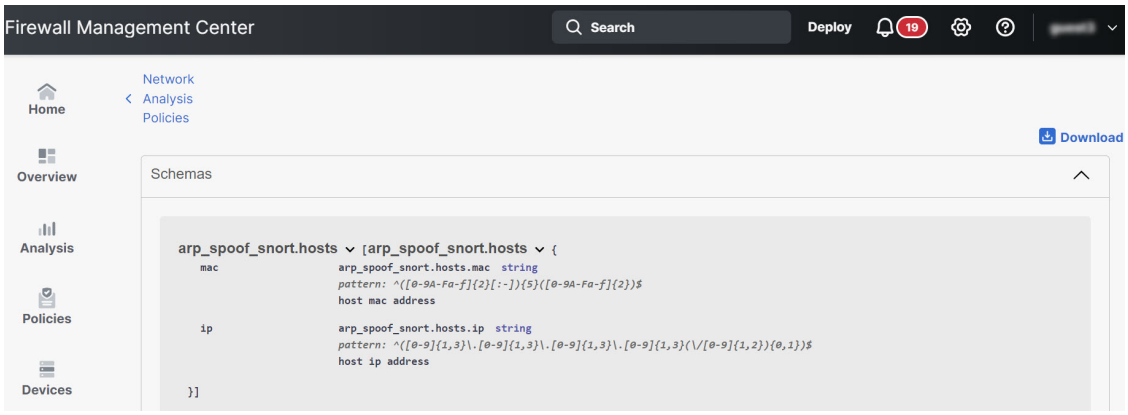

You must provide the values, range, patterns according to the accepted ones in the schema file to be able to successfully override the inspector configuration, otherwise, you get an error message.

b) Click **Download Sample File / Template** to use a pre-existing template that contains example configurations to help you with configuring the inspectors.

You can refer to the example configurations included in the sample file and make any changes that you may require.

c) Click **Download Full Configuration** to download the entire inspector configurations in a single JSON file.

Instead of expanding the inspectors separately, you can download the full configuration to look out for the information you need. All information regarding the inspector configuration is available in this file.

d) Click **Download Overridden Configuration** to download the inspector configuration that has been overridden.

Step 6 To override the existing configuration, follow the steps.

You can choose to override an inspector configuration using the following ways.

- Make inline edits for an inspector directly on the Firewall Management Center. See the topic **Make Inline Edit for an Inspector to Override Configuration** in the **Getting Started with Network Analysis Policies** chapter of the *Cisco Secure Firewall Management Center Snort 3 Configuration Guide*.
- Continue to follow the current procedure of using the **Actions** drop-down menu to upload the overridden configuration file.

If you chose to make inline edits directly in the Firewall Management Center, then you don't need to follow the current procedure further. Otherwise, you must follow this procedure completely.

a) Under **Inspectors**, expand the required inspector for which you want to override the default configuration.

The default configuration is displayed on the left column and the overridden configuration is displayed on the right column under the inspector.

You may need to search for an inspector by entering any relevant text in in the search bar.

b) Click the **Copy to clipboard** icon to copy the default inspector configuration to the clipboard.
c) Create a JSON file and paste the default configuration in it.
d) Keep the inspector configuration that you want to override, and remove all the other configuration and instances from the JSON file.

You can also use the **Sample File / Template** to understand how to override the default configuration. This is a sample file that includes JSON snippets explaining how you can customize the network analysis policy for Snort 3.

e) Make changes to the inspector configuration as needed.

Validate the changes and make sure they conform to the schema file. For multiton inspectors, make sure that the binder conditions for all instances are included in the JSON file. See *Multiton Inspectors* in the topic **Custom Network Analysis Policy Creation for Snort 3** in the *Cisco Secure Firewall Management Center Snort 3 Configuration Guide* for more information.

f) If you are copying any further default inspector configurations, append that inspector configuration to the existing file that contains the overridden configuration.

Note
The copied inspector configuration must comply with the JSON standards.

g) Save the overridden configuration file to your system.

Step 7 From the **Actions** drop-down menu, choose Upload Overridden Configuration to upload the JSON file that contains the overridden configuration.

Caution
Upload only the changes that you require. You should not upload the entire configuration as it makes the overrides sticky in nature and therefore, any subsequent changes to the default configuration as part of the LSP updates would not be applied.

You can drag and drop a file or click to browse to the JSON file saved in your system that contains the overridden inspector configuration.

- **Merge inspector overrides** – Content in the uploaded file is merged with the existing configuration if there is no common inspector. If there are common inspectors, then the content in the uploaded file (for common inspectors) takes precedence over the previous content, and it replaces the previous configuration for those inspectors.
- **Replace inspector overrides** – Removes all previous overrides and replaces them with the new content in the uploaded file.

Attention
Choosing this option deletes all the previous overrides. Make an informed decision before you override the configuration using this option.

If any error occurs while uploading the overridden inspectors, you see the error in the **Upload Overridden Configuration File** pop-up window. You can also download the file with the error, fix the error, and reupload the file.

Step 8 In the **Upload Overridden Configuration File** pop-up window, click **Import** to upload the overridden inspector configuration.

After you upload the overridden inspector configuration, you will see an orange icon next to the inspector that signifies that it is an overridden inspector.

Also, the **Overridden Configuration** column under the inspector shows the overridden value.

You can also view all the overridden inspectors using the **Show Overrides Only** checkbox adjacent to the Search bar.

Note

Make sure that you always download the overridden configuration, open the JSON file, and append any new changes/overrides to the inspector configurations to this file. This action is needed so that you do not lose the old overridden configurations.

Step 9 (Optional) Take a backup of the overridden configuration file on your system before making any new inspector configuration changes.

Tip
We recommend that you take the backup from time to time as you override the inspector configuration.

Related Topics

Revert Overridden Configuration to Default Configuration, on page 2131
View the List of Inspectors with Overrides, on page 2130
Search for an Inspector on the Network Analysis Policy Page, on page 2124
Copy the Inspector Configuration , on page 2125

Make Inline Edit for an Inspector to Override Configuration

For the Snort 3 version of the network analysis policy, you can make an inline edit for the inspector configuration to override the configuration according to your requirements.

Alternatively, you can also use the **Actions** drop-down menu to upload the overridden configuration file. See Customize the Network Analysis Policy, on page 2125 for more information.

Procedure

Step 1 Go to **Policies** > **Intrusion** > **Network Analysis Policies**.

Step 2 Go to the **Snort 3 Version** of the network analysis policy.

Step 3 Under **Inspectors**, expand the required inspector for which you want to override the default setting.

The default configuration is displayed on the left column and the overridden configuration is displayed on the right column under the inspector.

Step 4 Under the **Overridden Configuration** in the right column, click **Edit Inspector** (Pencil) icon to make changes to the inspector configuration.

The Override Configuration pop-up appears where you can make the required edits.

Note
- Make sure that you keep only those settings that you want to override. If you leave a setting with the same value, that field becomes sticky. This means if that setting is changed in the future by Talos, the current value will be retained.
- If you are adding or deleting any custom instance, make sure that you add or delete a binder rule for that instance in the binder inspector as well.

Step 5 Click **OK**.

If there are any errors according to the JSON standards, it shows you an error message.

Step 6 Click **Save** to save the changes.

If the changes conform to the OpenAPI schema specification, the Firewall Management Center allows you to save the configuration, otherwise, the **Error saving overridden configuration** pop-up appears that shows the errors. You can also download the file with the errors.

Related Topics

Customize the Network Analysis Policy, on page 2125
Revert Unsaved Changes during Inline Edits , on page 2130
Revert Overridden Configuration to Default Configuration, on page 2131
Examples of Custom Network Analysis Policy Configuration, on page 2134

Revert Unsaved Changes during Inline Edits

While making inline edits to override the configuration for an inspector , you can revert any unsaved changes. Note that this action reverts all unsaved changes to the most recently saved value, but does not revert the configuration to the default configuration for an inspector.

Procedure

Step 1 Go to **Policies** > **Intrusion** > **Network Analysis Policies**.

Step 2 Go to the **Snort 3 Version** of the network analysis policy.

Step 3 Under **Inspectors**, expand the required inspector for which you want to revert the unsaved changes.

The default configuration is displayed in the left column and the overridden configuration is displayed in the right column under the inspector.

Step 4 Under the **Overridden Configuration** on the right column, click the **Cross** ✕ icon to revert any unsaved changes for the inspector.

Alternatively, you can click **Cancel** to cancel the changes.

If you do not have any unsaved changes to the inspector configuration, then this option is not visible.

Related Topics

Revert Overridden Configuration to Default Configuration, on page 2131
Make Inline Edit for an Inspector to Override Configuration, on page 2129

View the List of Inspectors with Overrides

You can view a list of all the overridden inspectors.

Procedure

Step 1 Go to **Policies** > **Intrusion** > **Network Analysis Policies**.

Step 2 Go to the **Snort 3 Version** of the network analysis policy.

Step 3 Check the **Show Overrides Only** checkbox adjacent to the Search bar to view the list of overridden inspectors.

All the overridden inspectors are shown with an orange icon next to their names to help you identify them.

Related Topics

Search for an Inspector on the Network Analysis Policy Page, on page 2124
Make Inline Edit for an Inspector to Override Configuration, on page 2129
Customize the Network Analysis Policy, on page 2125

Revert Overridden Configuration to Default Configuration

You can revert any changes that you made to override the default configuration for an inspector. This action reverts the overridden configuration to the default configuration for an inspector.

Procedure

Step 1 Go to **Policies** > **Intrusion** > **Network Analysis Policies**.

Step 2 Go to the **Snort 3 Version** of the network analysis policy.

Step 3 Under **Inspectors**, expand the required inspector for which you want to revert the overridden configuration.

The overridden inspectors are shown with the orange icon next to their name.

The default configuration is displayed on the left column and the overridden configuration is displayed on the right column under the inspector. Under the **Overridden Configuration** on the right column, click **Revert to default configuration** (back arrow) icon to revert the overridden configuration for the inspector to the default configuration.

If you did not make any changes to the default configuration for the inspector, then this option is disabled.

Step 4 Click **Revert** to confirm the decision.

Step 5 Click **Save** to save the changes.

If you do not want to save the changes, you can click **Cancel** or the **Cross** ✕ icon.

Related Topics

Revert Unsaved Changes during Inline Edits , on page 2130
Customize the Network Analysis Policy, on page 2125
Make Inline Edit for an Inspector to Override Configuration, on page 2129
Examples of Custom Network Analysis Policy Configuration, on page 2134

Validate Snort 3 Policies

To validate the Snort 3 policies, here is a list of basic information that user can make note of:

- Current version of the Firewall Management Center can manage multiple Firewall Threat Defense versions.

- Current version of Firewall Management Center supports NAP configurations which are not applicable to previous version of Firewall Threat Defense devices.
- Current NAP Policy and validations will work based on the current version support.
- Changes may include content which is not valid for previous versions of Firewall Threat Defenses.
- Policy configuration changes are accepted if they are valid configuration for the current version and which is performed using current Snort 3 binary and NAP schema.
- For previous version Firewall Threat Defenses, validation is performed during deployment using NAP schema and Snort 3 binary for that specific version. If there is any configuration which is not applicable for the given version, user is provided information or warning that we will not deploy the configuration which is not supported on the given version and remaining configuration will get deployed.

In this procedure, when we associate the NAP policy to an Access Control Policy and deploy it on a device, for example any inspector like rate filter configuration is applied to validate the Snort 3 policies.

Procedure

Step 1 **Steps to Override NAP Policy Configuration:** Under **Inspectors** in the **Snort 3 Version** of the network analysis policy, expand the required inspector for which you want to override the default setting.

The default configuration is displayed on the left column and the overridden configuration is displayed on the right column under the inspector.

Step 2 Under the **Overridden Configuration** on the right column, click **Edit Inspector** (Pencil) icon to make changes to any inspector like rate_filter.

The Override Configuration pop-up appears where you can make the required edits to the rate_filter inspector.

Step 3 Click **OK**.

Step 4 Click **Save** to save the changes.

Alternatively, you can also use the **Actions** drop-down menu to upload the overridden configuration file.

Step 5 Click the **Actions** drop-down menu in the **Snort 3 Version** of the network analysis policy.

Step 6 Under **Upload** you can click **Overridden Configuration** to upload the JSON file that contains the overridden configuration.

Caution
Upload only the changes that you require. You should not upload the entire configuration as it makes the overrides sticky in nature and therefore, any subsequent changes to the default configuration as part of the LSP updates will not be applied.

You can drag and drop a file or click to browse to the JSON file saved in your system that contains the overridden inspector configuration.

- **Merge inspector overrides** – Content in the uploaded file is merged with the existing configuration if there is no common inspector. If there are common inspectors, then the content in the uploaded file (for common inspectors) takes precedence over the previous content, and it replaces the previous configuration for those inspectors.
- **Replace inspector overrides** – Removes all previous overrides and replaces them with the new content in the uploaded file.

Attention
As choosing this option deletes all the previous overrides, make an informed decision before you override the configuration using this option.

If any error occurs while uploading the overridden inspectors, you see the error on the **Upload Overridden Configuration File** pop-up window. You can also download the file with the error, then fix the error and reupload the file.

Step 7 **Steps to Associate NAP Policy to Access Control Policy:** In the access control policy editor, click **Advanced**, then click **Edit** next to the Network Analysis and Intrusion Policies section.

Step 8 From the **Default Network Analysis Policy** drop-down list, select a default network analysis policy.

If you choose a user-created policy, you can click **Edit** to edit the policy in a new window. You cannot edit system-provided policies.

Step 9 Click **OK**.

Step 10 Click **Save** to save the policy.

Step 11 Alternatively, in the access control policy editor, click **Advanced**, then click **Edit** next to the Network Analysis and Intrusion Policies section.

Step 12 Click **Add Rule**.

Step 13 Configure the rule's conditions by clicking the conditions you want to add.

Step 14 Click **Network Analysis** and choose the **Network Analysis Policy** you want to use to preprocess the traffic matching this rule.

Step 15 Click **Add**.

Step 16 **Deployment:** On the Firewall Management Center menu bar, click **Deploy** and then select **Deployment**.

Step 17 Identify and choose the devices on which you want to deploy configuration changes.

- Search—Search for the device name, type, domain, group, or status in the search box.
- Expand—Click **Expand Arrow** to view device-specific configuration changes to be deployed.

  By selecting the device check box, all the changes for the device, which are listed under the device, are pushed for deployment. However, you can use the **Policy Selection** to select individual policies or configurations to deploy while withholding the remaining changes without deploying them.

  Optionally, use **Show or Hide Policy** to selectively view or hide the associated unmodified policies.

Step 18 Click **Deploy**.

Step 19 If the system identifies errors or warnings in the changes to be deployed, it displays them in the **Validation Messages** window. To view complete details, click the arrow icon before the warnings or errors.

Note
It shows a warning that Snort 3 Network analysis policy contains inspectors or attributes that are not valid for this Firewall Threat Defense version, following the invalid settings will be skipped in deployment: Invalid inspectors are : [“rate_filter”] only for devices lower than 7.1 version.

Examples of Custom Network Analysis Policy Configuration

This is a sample file that includes JSON snippets explaining how you can customize the network analysis policy for Snort 3. You can choose to override an inspector configuration using the following ways:

- Make inline edits for an inspector directly on the Firewall Management Center. See Make Inline Edit for an Inspector to Override Configuration, on page 2129.
- Use the **Actions** drop-down menu to upload the overridden configuration file. See Customize the Network Analysis Policy, on page 2125.

Before you choose any of these options, review all the following details and examples that will help you in defining the network analysis policy overrides successfully. You must read and understand the examples for various scenarios explained here to avoid any risks and errors.

If you choose to override an inspector configuration from the **Actions** drop-down menu, you need to construct a JSON file for the network analysis policy overrides and upload the file.

For overriding an inspector configuration in the network analysis policy, you must upload only the changes that you require. You should not upload the entire configuration because it makes the overrides sticky in nature and therefore, any subsequent changes to the default values or configuration as part of the LSP updates would not be applied.

Here are the examples for various scenarios:

Enabling a Singleton Inspector when the Default State in the Base Policy is Disabled

```
{
  "rate_filter": {
    "enabled": true,
    "type": "singleton",
    "data": []
  }
}
```

Disabling a Singleton Inspector when the Default State in the Base Policy is Enabled

```
{
  "rate_filter": {
    "enabled": false,
    "type": "singleton",
    "data": []
  }
}
```

Enabling a Multiton Inspector when the Default State in the Base Policy is Disabled

```
{
  "ssh": {
    "enabled": true,
    "type": "multiton",
    "instances": []
  }
}
```

Disabling a Multiton Inspector when the Default State in the Base Policy is Enabled

```
{
  "ssh": {
```

```
    "enabled": false,
    "type": "multiton",
    "instances": []
  },
  "iec104": {
    "type": "multiton",
    "enabled": false,
    "instances": []
  }
}
```

Overriding the Default Value of Specific Setting(s) for Singleton Inspector

```
{
  "normalizer": {
    "enabled": true,
    "type": "singleton",
    "data": {
      "tcp": {
        "block": true
      },
      "ip6": true
    }
  }
}
```

Overriding Specific Setting(s) of a Default Instance (where Instance Name Matches with Inspector Type) in Multiton Inspector

```
{
  "http_inspect": {
    "enabled": true,
    "type": "multiton",
    "instances": [
      {
        "data": {
          "unzip": false
        },
        "name": "http_inspect"
      }
    ]
  }
}
```

Adding Binder Rule for a Default Instance with Required Changes

Note Default binder rules can't be edited, they are always appended at the end.

```
{
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "use": {
          "type": "http_inspect"
        },
        "when": {
          "role": "server",
```

```
          "service": "http",
          "dst_nets": "10.1.1.0/24"
        }
      }
    ]
  }
}
```

Adding a New Custom Instance

Note Corresponding binder rule entry must be defined in the binder inspector.

```
{
  "telnet": {
    "enabled": true,
    "type": "multiton",
    "instances": [
      {
        "name": "telnet_my_instance",
        "data": {
          "encrypted_traffic": true
        }
      }
    ]
  },
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "when": {
          "role": "any",
          "service": "telnet"
        },
        "use": {
          "type": "telnet",
          "name": "telnet_my_instance"
        }
      }
    ]
  }
}
```

Overriding a Singleton Instance, Multiton Default Instance, and Creating a New Multiton Instance in a Single JSON Override

Example to show the following in a single JSON override:

- Overriding a Singleton instance (**`normalizer`** inspector)
- Overriding a Multiton default instance (**`http_inspect`** inspector)
- Creating a new Multiton instance (**`telnet`** inspector)

```
{
  "normalizer": {
    "enabled": true,
    "type": "singleton",
    "data": {
```

```
      "tcp": {
        "block": true
      },
      "ip6": true
    }
  },
  "http_inspect": {
    "enabled": true,
    "type": "multiton",
    "instances": [
      {
        "data": {
          "unzip": false,
          "xff_headers": "x-forwarded-for true-client-ip x-another-forwarding-header"
        },
        "name": "http_inspect"
      }
    ]
  },
  "telnet": {
    "enabled": true,
    "type": "multiton",
    "instances": [
      {
        "name": "telnet_my_instance",
        "data": {
          "encrypted_traffic": true
        }
      }
    ]
  },
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "when": {
          "role": "any",
          "service": "telnet"
        },
        "use": {
          "type": "telnet",
          "name": "telnet_my_instance"
        }
      },
      {
        "use": {
          "type": "http_inspect"
        },
        "when": {
          "role": "server",
          "service": "http",
          "dst_nets": "10.1.1.0/24"
        }
      }
    ]
  }
}
```

Note You don't need to give the **name** attribute for the default instance in binder rules.

Configuring arp_spoof

Example for configuring **arp_spoof**:

Ther **arp_spoof** inspector does not have any default configurations for any attributes. This demonstrates the case where you can provide the overrides.

```
{
  "arp_spoof": {
    "type": "singleton",
    "data": {
      "hosts": [
        {
          "ip": "1.1.1.1",
          "mac": "ff:0f:f1:0f:0f:ff"
        },
        {
          "ip": "2.2.2.2",
          "mac": "ff:0f:f2:0f:0f:ff"
        }
      ]
    },
    "enabled": true
  }
}
```

Configuring rate_filter

```
{
  "rate_filter": {
    "data": [
      {
        "apply_to": "[10.1.2.100, 10.1.2.101]",
        "count": 5,
        "gid": 135,
        "new_action": "alert",
        "seconds": 1,
        "sid": 1,
        "timeout": 5,
        "track": "by_src"
      }
    ],
    "enabled": true,
    "type": "singleton"
  }
}
```

Configuring Binder Rules when Multi-Hierarchy Network Analysis Policy is Used

This example illustrates adding a new custom instance in child policy and the way binder rules should be written. Binder rules are defined as a list and therefore, it is important to pick up the rules defined in the parent policy and build the new rules on top of it as rules will not be merged automatically. The binder rules available in child policy are a source of truth in totality.

On the Firewall Threat Defense, the default Cisco Talos policy rules are appended on these user-defined overrides.

Parent Policy:

We have defined a custom instance by the name **telnet_parent_instance** and the corresponding binder rule.

```
{
  "telnet": {
    "type": "multiton",
    "instances": [
      {
        "data": {
          "normalize": true,
          "encrypted_traffic": true
        },
        "name": "telnet_parent_instance"
      }
    ],
    "enabled": true
  },
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "when": {
          "role": "any",
          "service": "telnet"
        },
        "use": {
          "type": "telnet",
          "name": "telnet_parent_instance"
        }
      }
    ]
  }
}
```

Child Policy:

This network analysis policy has the aforementioned policy as its base policy. We have defined a custom instance by the name **`telnet_child_instance`** and have also defined the binder rules for this instance. The binder rules from parent policy need to be copied here, and then child policy binder rules can be prepended or appended on top of it based on the nature of the rule.

```
{
  "telnet": {
    "type": "multiton",
    "instances": [
      {
        "data": {
          "normalize": true,
          "encrypted_traffic": false
        },
        "name": "telnet_child_instance"
      }
    ],
    "enabled": true
  },
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "when": {
          "role": "any",
          "service": "telnet",
          "nets": "10.2.2.0/24"
        },
        "use": {
```

```
          "type": "telnet",
          "name": "telnet_child_instance"
        }
      },
      {
        "when": {
          "role": "any",
          "service": "telnet"
        },
        "use": {
          "type": "telnet",
          "name": "telnet_parent_instance"
        }
      }
    ]
  }
}
```

Configuring List Inspector Attribute in General

While changing overrides for any attribute of type list, it is important to pass the full contents rather than partial override. This means if a base policy attributes are defined as:

```
{
  "list-attribute": [
    {
      "entry1": {
        "key1": "value1"
      }
    },
    {
      "entry2": {
        "key2": "value2"
      }
    }
  ]
}
```

If you want to modify **`value1`** to **`value1-new`**, the override payload must look like the following:

Correct Way:

```
{
  "list-attribute": [
    {
      "entry1": {
        "key1": "value1-new"
      }
    },
    {
      "entry2": {
        "key2": "value2"
      }
    }
  ]
}
```

Incorrect Way:

```
{
  "list-attribute": [
    {
      "entry1": {
        "key1": "value1-new"
      }
```

```
    }
  ]
}
```

You can understand this configuration by taking the trimmed values of the **alt_max_command_line_len** attribute in the **smtp** inspector. Suppose the default (base) policy configuration for **smtp** inspector is as follows:

```
{
  "smtp": {
    "type": "multiton",
    "instances": [
      {
        "name": "smtp",
        "data": {
          "decompress_zip": false,
          "normalize_cmds": "ATRN AUTH BDAT CHUNKING DATA DEBUG EHLO
           EMAL ESAM ESND ESOM ETRN EVFY EXPN HELO HELP IDENT MAIL
           NOOP ONEX QUEU QUIT RCPT RSET SAML SEND SOML STARTTLS TICK
           TIME TURN TURNME VERB VRFY X-ADAT XADR XAUTH XCIR X-DRCP X-
           ERCP XEXCH50 X-EXCH50 X-EXPS XGEN XLICENSE X-LINK2STATE XQUE
           XSTA XTRN XUSR",
          "ignore_data": false,
          "max_command_line_len": 512,
          "max_header_line_len": 1000,
          "log_rcptto": false,
          "decompress_swf": false,
          "max_response_line_len": 512,
          "b64_decode_depth": -1,
          "max_auth_command_line_len": 1000,
          "log_email_hdrs": false,
          "xlink2state": "alert",
          "binary_data_cmds": "BDAT XEXCH50",
          "auth_cmds": "AUTH XAUTH X-EXPS",
          "log_filename": false,
          "uu_decode_depth": -1,
          "ignore_tls_data": false,
          "data_cmds": "DATA",
          "bitenc_decode_depth": -1,
          "alt_max_command_line_len": [
            {
              "length": 255,
              "command": "ATRN"
            },
            {
              "command": "AUTH",
              "length": 246
            },
            {
              "length": 255,
              "command": "BDAT"
            },
            {
              "length": 246,
              "command": "DATA"
            }
          ],
          "log_mailfrom": false,
          "decompress_pdf": false,
          "normalize": "none",
          "email_hdrs_log_depth": 1464,
          "valid_cmds": "ATRN AUTH BDAT CHUNKING DATA DEBUG EHLO
           EMAL ESAM ESND ESOM ETRN EVFY EXPN HELO HELP IDENT MAIL
           NOOP ONEX QUEU QUIT RCPT RSET SAML SEND SOML STARTTLS TICK
```

```
            TIME TURN TURNME VERB VRFY X-ADAT XADR XAUTH XCIR X-DRCP X-
            ERCP XEXCH50 X-EXCH50 X-EXPS XGEN XLICENSE X-LINK2STATE XQUE
            XSTA XTRN XUSR",
          "qp_decode_depth": -1
        }
      }
    ],
    "enabled": true
  }
}
```

Now, if you want to add two more objects to the **alt_max_command_line_len** list:

```
{
    "length": 246,
    "command": "XEXCH50"
},
{
    "length": 246,
    "command": "X-EXPS"
}
```

Then the custom network analysis policy override JSON would look like the following:

```
{
  "smtp": {
    "type": "multiton",
    "instances": [
      {
        "name": "smtp",
        "data": {
          "alt_max_command_line_len": [
            {
              "length": 255,
              "command": "ATRN"
            },
            {
              "command": "AUTH",
              "length": 246
            },
            {
              "length": 255,
              "command": "BDAT"
            },
            {
              "length": 246,
              "command": "DATA"
            },
            {
              "length": 246,
              "command": "XEXCH50"
            },
            {
              "length": 246,
              "command": "X-EXPS"
            }
          ]
        }
      }
    ],
    "enabled": true
  }
}
```

Configuring Overrides when Multi-Hierarchy Network Analysis Policy is used in Multiton Inspector

This example illustrates overriding attributes in child policy and how the merged configuration will be used in the child policy for any instance. Any overrides defined in the child policy will be merged with the parent policy. Thus, if attribute1 and attribute2 are overridden in parent policy and attribute2 and attribute3 are overridden in the child policy, the merged configurations are for child policy. This means that attribute1 (defined in parent policy), attribute2 (defined in child policy), and attribute3 (defined in child policy) will be configured on the device.

Parent Policy:

Here we have defined a custom instance by the name **`telnet_parent_instance`** and overridden 2 attributes namely, **`normalize`** and **`encrypted_traffic`** in the custom instance.

```
{
  "telnet": {
    "type": "multiton",
    "instances": [
      {
        "data": {
          "normalize": true,
          "encrypted_traffic": false
        },
        "name": "telnet_parent_instance"
      }
    ],
    "enabled": true
  },
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "when": {
          "role": "any",
          "service": "telnet"
        },
        "use": {
          "type": "telnet",
          "name": "telnet_parent_instance"
        }
      }
    ]
  }
}
```

Child Policy:

This network analysis policy has the aforementioned policy as its base policy. We have overridden attribute **`encrypted_traffic`** from parent policy and also overridden new attribute **`ayt_attack_thresh`**.

```
{
  "telnet": {
    "type": "multiton",
    "instances": [
      {
        "data": {
          "encrypted_traffic": true,
          "ayt_attack_thresh": 1
        },
        "name": "telnet_parent_instance"
      }
    ],
    "enabled": true
```

```
  }
}
```

With the above policy JSON, when you deploy the network analysis policy the following merged JSON will be configured on the device.

```
{
  "telnet": {
    "type": "multiton",
    "instances": [
      {
        "data": {
          "normalize": true,
          "encrypted_traffic": true,
          "ayt_attack_thresh": 1
        },
        "name": "telnet_parent_instance"
      }
    ],
    "enabled": true
  },
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "when": {
          "role": "any",
          "service": "telnet"
        },
        "use": {
          "type": "telnet",
          "name": "telnet_parent_instance"
        }
      }
    ]
  }
}
```

This example illustrates details for the custom network analysis policy. The same behavior is also exhibited in the default instance. Also, a similar merging would be done for Singleton inspectors.

Removing all the Inspector Overrides for the Network Analysis Policy:

Whenever you want to remove all the overrides for a specific network analysis policy, you can upload an empty JSON. While uploading the overrides, choose the option **Replace inspector overrides**.

```
{
}
```

Related Topics

Snort 3 Definitions and Terminologies for Network Analysis Policy, on page 2115
Network Analysis Policy Mapping, on page 2122
Custom Network Analysis Policy Creation for Snort 3, on page 2117
Search for an Inspector on the Network Analysis Policy Page, on page 2124
Copy the Inspector Configuration , on page 2125
Customize the Network Analysis Policy, on page 2125
View the List of Inspectors with Overrides, on page 2130

Network Analysis Policy Settings and Cached Changes

When you create a new network analysis policy, it has the same settings as its base policy.

When tailoring a network analysis policy, especially when disabling inspectors, keep in mind that some inspectors and intrusion rules require that traffic first be decoded or preprocessed in a certain way. If you disable a required inspector, the system automatically uses it with its current settings, although the inspector remains disabled in the network analysis policy web interface.

Note Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet **must** complement each other. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task.

The system caches one network analysis policy per user. While editing a network analysis policy, if you select any menu or other path to another page, your changes stay in the system cache even if you leave the page.

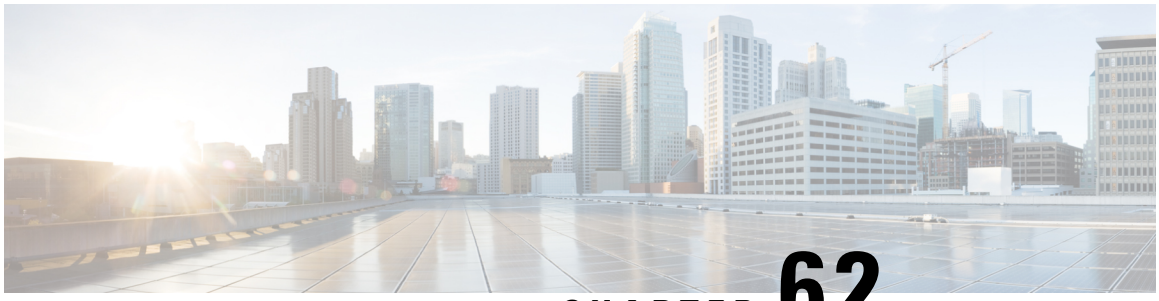

CHAPTER 62

Encrypted Visibility Engine

Encrypted Visibility Engine (EVE) is used to identify client applications and processes utilizing TLS encryption. It enables visibility and allows administrators to take actions and enforce policy within their environments. The EVE technology can also be used to identify and stop malware.

- Overview of Encrypted Visibility Engine, on page 2147
- How EVE Works, on page 2148
- Indications of Compromise Events, on page 2149
- QUIC Fingerprinting in EVE, on page 2149
- Configure EVE, on page 2149
- Configure EVE Exception Rules, on page 2152
- Event Enrichment, on page 2154

Overview of Encrypted Visibility Engine

The encrypted visibility engine (EVE) is used to provide more visibility into the encrypted sessions without the need to decrypt them. These insights into encrypted sessions are obtained by Cisco's open-source library that is packaged in Cisco's vulnerability database (VDB). The library fingerprints and analyzes incoming encrypted sessions and matches it against a set of known fingerprints. This database of known fingerprints is also available in the Cisco VDB.

Note The encrypted visibility engine feature is supported only on Firewall Management Center-managed devices running Snort 3. This feature is not supported on Snort 2 devices and Firewall Device Manager-managed devices.

Some of the important features of EVE are the following:

- You can take access control policy actions on the traffic using information derived from EVE.
- The VDB included in Cisco Secure Firewall has the ability to assign applications to some processes detected by EVE with a high confidence value. Alternatively, you can create custom application detectors to:
  - Map EVE-detected processes to new user-defined applications.
  - Override the built-in value of process confidence that is used to assign applications to EVE-detected processes.

See the **Configuring Custom Application Detectors** and **Specifying EVE Process Assignments** sections in the **Application Detection** chapter of the Cisco Secure Firewall Management Center Device Configuration Guide.

- EVE can detect the operating system type and version of the client that created a Client Hello packet in the encrypted traffic.
- EVE supports fingerprinting and analysis of Quick UDP Internet Connections (QUIC) traffic too. The server name from the Client Hello packet is displayed in the URL field of the **Connection Events** page.

Attention To use EVE on Firewall Management Center, you must have a valid IPS license on your device. In the absence of a IPS license, the policy displays a warning and deployment is not allowed.

Note

- EVE can detect the operating system type and version of SSL sessions. Normal usage of the operating system, such as running applications and package management software, can trigger OS detection. To view client OS detection, in addition to enabling the EVE toggle button, you must enable **Hosts** under **Policies** > **Network Discovery**. To view a list of possible operating systems on the host IP address, click **Analysis** > **Hosts** > **Network Map**, and then choose the required host.
- After enabling EVE for your access control policy, ensure that you have turned on logging for the access control rules within that policy to display the expected results on the EVE dashboard whenever any specific rule conditions are met. For more information on how to turn on logging, see Create and Edit Access Control Rules.
- EVE will not provide visibility or insights for encapsulated traffic.

Related Links

Configure EVE, on page 2149

Configure EVE Exception Rules, on page 2152

How EVE Works

The Encrypted Visibility Engine (EVE) inspects the Client Hello portion of the TLS handshake to identify client processes. The Client Hello is the initial data packet that is sent to the server. This gives a good indication of the client process on the host. This fingerprint, combined with other data such as destination IP address, provides the basis for EVE's application identification. By identifying specific application fingerprints in the TLS session establishment, the system can identify the client process and take appropriate action (allow/block).

EVE can identify over 5,000 client processes. The system maps a number of these processes to client applications for use as criteria in access control rules. This gives the system the ability to identify and control these applications without enabling TLS decryption. By using fingerprints of known malicious processes, EVE technology can also be used to identify and block encrypted malicious traffic without outbound decryption.

Through machine learning (ML) technology, Cisco processes over one billion TLS fingerprints and over 10000 malware samples daily to create and update EVE fingerprints. These updates are then delivered to customers using Cisco Vulnerability Database (VDB) package.

If EVE does not recognize a fingerprint, it identifies client application and estimates the threat score of the first flow using the destination details, such as IP address, port, and server name. At this point, the status of the fingerprints are randomized and the status can be viewed in the debug logs. For subsequent flows with the same fingerprint, EVE skips reanalysis and marks the fingerprint status as unlabeled. If you intend to block traffic based on EVE's Low or Very Low score thresholds, the initial flow is blocked. However, future flows will be allowed once the application's fingerprint is cached.

Indications of Compromise Events

The host's Indications of Compromise (IoC) events for encrypted visibility engine detection allows you to check connection events with a very high malware confidence level, as reported by EVE. IoC events are triggered for encrypted sessions generated from a host using a malicious client. You can view information, such as IP address, MAC address, and OS information of the malicious host, and the timestamp of the suspicious activity.

A session with Encrypted Visibility Threat Confidence score 'Very High' as seen in connection events genreates an IoC event. You must enable **Hosts** from **Policies** > **Network Discovery**. In the Firewall Management Center, you can view the IoC event existence from:

- **Analysis** > **Indications of Compromise**.
- **Analysis** > **Network Map** > **Indications of Compromise** > Choose the host that must be checked.

  You can view the process information of the session that generated the IoC on the **Connection Events** page. Click **Analysis** > **Connections Header** > **Events** to access the **Connection Events** page. Note that you must manually select the Encrypted Visibility fields and IoC field from the **Table View of Connection Events** tab.

QUIC Fingerprinting in EVE

Snort can identify client applications in Quick UDP Internet Connections (QUIC sessions) based on EVE. QUIC fingerprinting can:

- Detect applications over QUIC without enabling decryption.
- Identify malware without enabling decryption.
- Detect service applications. You can assign access control rules based on the service detected over the QUIC protocol.

Configure EVE

Procedure

Step 1 Choose **Policies** > **Access Control**.

Step 2 Click **Edit** (✎) next to the access control policy you want to edit.

Step 3 Choose **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 4 Click **Edit** (✎) next to **Encrypted Visibility Engine (EVE)**.

Step 5 On the **Encrypted Visibility Engine** page, enable the **Encrypted Visibility Engine (EVE)** toggle button.

Step 6 **Use EVE for Application Detection**—This toggle button is enabled by default, which means that EVE is allowed to assign client applications to processes.

EVE's fingerprint information is added in the **Encrypted Visibility Fingerprint** column header of the connection events or unified events. For further analysis of the EVE data collected, you can right-click the fingerprint information to open a dropdown menu. In the menu, click **View Encrypted Visibility Engine Process Analysis** to go to the Secure Firewall Application Detectors site and view details, such as the fingerprint, VDB version, and so on. Different rows with the same fingerprint string and potential process names associated with them and their prevalence are displayed. Prevalence indicates the frequency of a process associated with a particular fingerprint in the data collection system. You can choose the process names and click **Submit Request** to give feedback about any discrepancy in EVE's process detection. For example, you can submit requests if the process name that is detected does not match with the traffic that is being sent or if the process name is not detected at all for a particular fingerprint.

Access to view additional details on the Secure Firewall Application Detectors page is currently not available to users with non-Cisco email addresses.

If you disable the **Use EVE for Application Detection** toggle button:

- AppID-identified clients are assigned to processes and you can see the EVE process and score, but there is no mapping of EVE-detected processes to applications and no action is taken. You can see the details of the events under **Connection Events** or **Unified Events**. To see the difference in connection events (with and without application assignment), see the **Client Application** column header.
- The **Encrypted Visibility Fingerprint** field in the connection events or unified events is empty.

Step 7 Enable the **Block Malware Processes Based on Threat Confidence Level** toggle button to block malicious client processes with the prefix *malware_* based on EVE's threat confidence level.

The default block threshold is 99 percent, which means:

- If EVE detects the traffic to be malware with 99 percent confidence or more, the traffic is blocked.
- If EVE detects the traffic to be malware with less than 99 percent confidence, EVE takes no action.

Note
If EVE has blocked the traffic, the **Reason** column header on the **Connection Events** page displays **Encrypted Visiblity Block**.

Step 8 Use the slider to adjust the threshold for blocking based on EVE's threat confidence, which ranges from **Very Low** to **Very High**.

Step 9 For further granular control, enable the **Advanced Mode** toggle button. Now, you can assign a specific EVE Threat Confidence Level for blocking traffic. The default block threshold is 99 percent.

Caution
We recommend that you do not set a threshold below 50 percent to ensure optimal performance.

Step 10 Click **OK**.

Step 11 Click **Save**.

What to do next

Deploy configuration changes.

View Encrypted Visibility Engine Events

After enabling the **Encrypted Visibility Engine** and deploying your access control policy, you can start sending live traffic through your system. You can view the logged connection events on the **Connection Events** page or on the **Unified Events** page.

Perform this procedure to access the connection events in the Firewall Management Center.

Procedure

Step 1 Click **Analysis** > **Connections Header** > **Events**.

Step 2 Click the **Table View of Connection Events** tab.

You can also view the connection events in the **Unified Events** page. Click **Analysis** > **Unified Events** to access the **Unified Events** page.

The Encrypted Visibility Engine can identify the client process that initiated a connection and the operating system in the client, and indicate if the process contains malware or not.

On the **Connection Events** page, you must explicitly enable these columns that are added for the Encrypted Visibility Engine:

- **EVE Process Name**
- **EVE Process Confidence Score**
- **EVE Threat Confidence**
- **EVE Threat Confidence Score**
- **Detection Type**

For information about these fields, see Connection and Security-Related Connection Event Fields in the Cisco Secure Firewall Management Center Administration Guide.

Note
On the **Connection Events** page, if processes are assigned applications, the **Detection Type** column displays **Encrypted Visibility Engine**, indicating that the client application was identified by the Encrypted Visibility Engine. Without application assignments to process names, the **Detection Type** column displays **AppID**, indicating that the engine that identified the client application was AppID.

View EVE Dashboard

You can view the EVE analysis information in the following dashboards:

Before you begin

- In an access control policy, the **Encrypted Visisibility Engine (EVE)** must be enabled under the **Advanced Settings**.
- To view the **Connections with Detected Process** and **Malicious Processes** widgets, the devices must be running Firewall Threat Defense Version 7.7 or later.

Procedure

Step 1 Go to **Overview** > **Dashboards**, and click **Dashboard**.

Step 2 In the **Summary Dashboard** window, click the **Encrypted Visibility Engine** tab.

Step 3 You can view the following dashboards:

- **Discovered Processes**—Displays top client processes used in your network and the connection count. You can click the process name in the table to see the filtered view of the **Connection Events** page, which is filtered by the process name.
- **Threat Confidence**—Displays connections by the confidence levels. You can click the Threat confidence level in the table to see the filtered view of the **Connection Events** page, which is filtered by the confidence level.
- **Connections with Detected Process**—Displays total count of connections in which EVE identified the client processes.
- **Malicious Processes**—Displays count of malicious client processes identified by EVE with high and very high threat confidence levels.

Configure EVE Exception Rules

You can create an encrypted visibility engine (EVE) exception rule to ensure the continuity of trusted connections and services by bypassing the EVE's block action. You can add attributes such as process names, source and destination IP addresses/ FQDNs, and dynamic objects to the exception rule. For example, you may want to bypass EVE's block verdict for trusted networks. All the connections in the bypassed networks are exempted from EVE's block verdict based on the threat confidence level.

Procedure

Step 1 Choose **Policies** > **Access Control**.

Step 2 Click **Edit** (✎) next to the access control policy you want to edit.

Step 3 From the **More** drop-down arrow at the end of the packet flow line, choose **Advanced Settings**.

Step 4 Next to **Encrypted Visibility Engine (EVE)**, click **Edit** (✎).

Step 5 On the **Encrypted Visibility Engine** page, click the **Encrypted Visibility Engine (EVE)** toggle button to enable EVE.

Step 6 Enable the **Block based on EVE threat confidence level** toggle button to block traffic based on EVE's threat confidence level.

Step 7 Click **Add Exception Rule** and add one or more of the following attributes.

a) Under the **Process Name** tab, enter an EVE-identified process name, and click **Add to Process** on the right side of the window.

You can add multiple process names to the same exception rule. EVE exception list based on process names works only with EVE-identified process names, which are case- and space-sensitive.

b) Under the **Network Objects** tab, perform one of the following:

- Choose one or more IP addresses or FQDNs from the list and click either **Add to Source Network** or **Add to Destination Network**.
- Under **Selected Source Network** or **Selected Destination Network**, manually enter the IP address, and click the **Add** (+) icon to add it to the list of selected networks.

c) Under the **Dynamic Attributes** tab, choose the dynamic objects and add it to the **Selected Dynamic Objects** list.

For more information about creating dynamic objects or working with dynamic objects, see the *Create Dynamic Objects for the First Time* or *Work With Dynamic Objects* sections in Secure Firewall Management Center Device Configuration Guide.

d) (Optional) In the **Comment** field available on all the tabs, you can enter a reason for adding the required attributes to the EVE exception rule.

Step 8 Click **Save** to save the EVE exception rule.

Step 9 Save and deploy the access control policy on the devices.

Note When a connection matches an exception rule, it bypasses the EVE's block verdict. You can view EVE's action in the **Connection Events** or **Unified Events** page. The **Reason** column header displays **EVE Exempted** for identification of such EVE-bypassed traffic.

Add Exception Rule from Unified Events

Use the **Unified Events** page to add exception rules for connections that are blocked by EVE.

Before you begin

Exception list is supported only from threat defense Version 7.6.0 or later.

Procedure

Step 1 Click **Analysis** > **Unified Events**.

Step 2 In the **Reason** column with **Encrypted Visibility Block** as the reason, click the **Ellipsis**( ) icon inside the cell.

Step 3 Choose **Add EVE Exception Rule** from the drop-down list.

Step 4 In the **Encrypted Visibility Engine** window that is displayed, the rule is automatically added to the bottom of the exception list. You can review and make changes to the added rule before saving and deploying the configuration.

Event Enrichment

Context enrichment for MITRE ATT&CK occurs from the Talos taxonomy and the encrypted visibility engine (EVE). Both Talos and EVE enrichments are communicated using the Talos taxonomy. EVE enrichment works when EVE is enabled. For more information about enabling EVE, see Configure EVE, on page 2149.

On the **Connection Events** page, you can view the following column headers that are added as part of enriched eventing content. You must explicitly enable these columns.

- **MITRE ATT&CK**
- **Other Enrichment**

For information about these fields, see *Connection and Security-Related Connection Event Fields* in Cisco Secure Firewall Management Center Administration Guide.

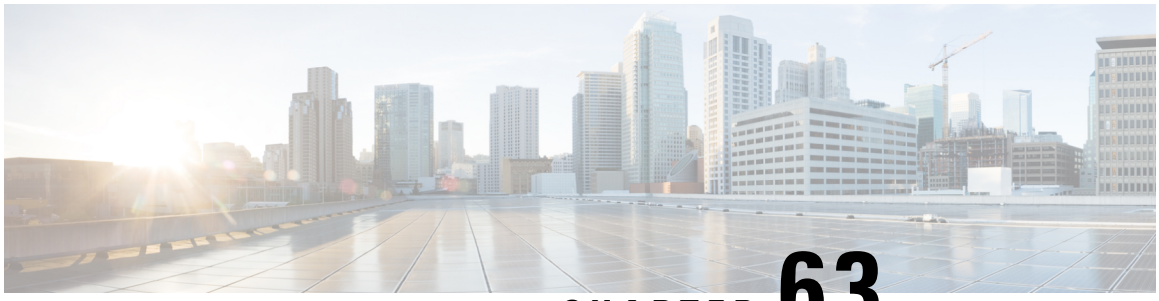

CHAPTER 63

Elephant Flow Detection

Elephant flows are extremely large (in total bytes), continuous flows set up by a TCP (or other protocols) flow measured over a network link. By default, elephant flows are those larger than 1 GB/10 seconds. They can cause performance duress in Snort cores. Elephant flows are not numerous, but they can occupy a disproportionate share of the total bandwidth over a period of time. They can lead to problems, such as high CPU utilization, packet drops, and so on.

From Firewall Management Center 7.2.0 onwards (Snort 3 devices only), you can use the elephant flow feature to detect and remediate elephant flows, which helps to reduce system stress and resolve the mentioned issues.

- About Elephant Flow Detection and Remediation, on page 2155
- Elephant Flow Upgrade from Intelligent Application Bypass, on page 2155
- Configure Elephant Flow, on page 2156

About Elephant Flow Detection and Remediation

You can use the elephant flow detection feature to detect and remediate elephant flows. The following remediation actions can be applied:

- **Bypass elephant flow**–You can configure elephant flow to bypass Snort inspection. If this is configured, Snort does not receive any packet from that flow.
- **Throttle elephant flow**–You can apply rate-limit to the flow and continue to inspect flows. The flow rate is calculated dynamically and 10% of the flow rate is reduced. Snort sends the verdict (QoS flow with 10% less flow rate) to the firewall engine. If you choose to bypass all applications including unidentified applications, you cannot configure the throttle action (rate-limit) for any flow.

Note For the elephant flow detection to work, Snort 3 must be the detection engine.

Elephant Flow Upgrade from Intelligent Application Bypass

Intelligent Application Bypass (IAB) is deprecated from version 7.2.0 onwards for Snort 3 devices.

For devices running 7.2.0 or later, you must configure elephant flow settings under the **Elephant Flow Settings** section in the AC policy (Advanced settings tab).

Post-upgrade to 7.2.0 (or later), if you are using a Snort 3 device, the elephant flow configuration settings will be picked and deployed from the **Elephant Flow Settings** section and not from the **Intelligent Application Bypass Settings** section, so if you have not migrated to Elephant Flow configuration settings, your device will lose the elephant flow configuration upon the next deployment.

The following table shows the IAB or elephant flow configurations that can be applied to version 7.2.0 or later and to version 7.1.0 or earlier that are running Snort 3 or Snort 2 engines.

Firewall Management Center	Firewall Threat Defense	Elephant Flow or IAB Configuration
Firewall Management Center 7.0 or 7.1	Snort 2 device	Configuration from IAB is applicable.
	Snort 3 device	Configuration from IAB is applicable.
Firewall Management Center 7.2.0	Snort 2 device	Configuration from IAB is applicable.
	Snort 3 device (7.1.0 and earlier)	Configuration from IAB is applicable.
	Snort 3 device (7.2.0 and later)	Configuration from Elephant Flow is applicable.

Configure Elephant Flow

You can configure elephant flow to take actions on elephant flows, which helps resolve issues, such as system duress, high CPU utilization, packet drops, and so on.

Attention Elephant flow detection is not applicable for prefiltered, trusted, or fast-forwarded flows, which do not process through Snort. As elephant flows are detected by Snort, elephant flow detection is not applicable for encrypted traffic.

Procedure

Step 1 In the access control policy editor, click **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line. Then, click **Edit** (✎) next to **Elephant Flow Settings**.

If **View** (👁) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Figure 447: Configure Elephant Flow Detection

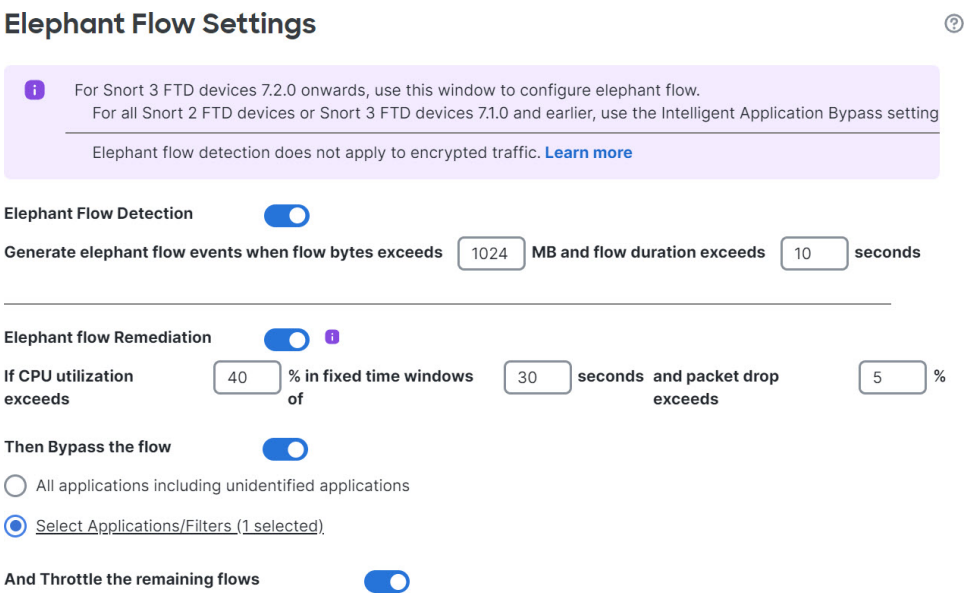

Figure 448: Configure Elephant Flow Detection

Step 2 The **Elephant Flow Detection** toggle button is enabled by default. You can configure the values for flow bytes and flow duration. When they exceed your configured values, elephant flow events are generated.

Step 3 To remediate elephant flows, enable the **Elephant Flow Remediation** toggle button.

Step 4 To set the criteria for remediation of the elephant flow, configure the values for CPU utilization %, duration of fixed time windows, and packet drop %.

CPU utilization is calculated per elephant flow and is derived from the flow latency. If the CPU utilization crosses the configured threshold and other configurations, such as fixed time windows and packet drops, are also matched, the elephant flow remediation actions are applied. Similarly, packet drop calculation is based on the packets dropped per CPU. After the packet drop percentage exceeds the configured value on a specific CPU, the remediation actions are applied. For example, consider that configurations are set to the default, that is, CPU utilization of 40%, fixed time window of 30 seconds, and packet drop of 5%. On a specific CPU, if more than 5% of packet drops are detected and the CPU utilization per flow exceeds 40% in the fixed time frame of 30 seconds, then the flows are either bypassed or throttled.

Step 5 You can perform the following actions for elephant flow remediation when it meets the configured criteria:

a. **Bypass the flow**—Enable this button to bypass Snort inspection for selected applications or filters. Choose from:

- **All applications including unidentified applications**—Select this option to bypass all the application traffic. If you configure this option, you cannot configure the throttle action (rate-limit) for any flow.
- **Select Applications/Filters**—Select this option to select the applications or filters whose traffic you want to bypass; see the topic **Configuring Application Conditions and Filters** in the **Access Control Rules** chapter in the Cisco Secure Firewall Management Center Device Configuration Guide.

b. **Throttle the flow**—Enable this button to apply rate-limit to the flow and continue to inspect flows. Note that you can select the applications or filters to bypass Snort inspection and throttle the remaining flows.

Note
Automatic removal of throttle from a throttled elephant flow occurs when the system is out of duress, that is, the percentage of Snort packet drops is lesser than your configured threshold. Consequently, rate limiting is also removed.

You can also manually remove throttling from a throttled elephant flow, using the following threat defense commands:

- **clear efd-throttle <5-tuple/all> bypass**—This command removes throttling from the throttled elephant flow and bypasses Snort inspection.
- **clear efd-throttle <5-tuple/all>**—This command removes throttling from the throttled elephant flow and Snort inspection continues. Elephant flow remediation is skipped after using this command.

  For more information about these commands, see the Cisco Secure Firewall Threat Defense Command Reference.

Step 6 In the **Remediation Exemption Rule** section, click **Add Rule** to configure L4 access control list (ACL) rules for flows that must be exempted from remediation.

Step 7 In the **Add Rule** window, use the **Networks** tab to add the network details, that is the source network and the destination network. Use the **Ports** tab to add the source port and the destination port.

If an elephant flow is detected and it matches the rules that are defined, an event is generated with the reason as **Elephant Flow Exempted** in the **Reason** column header of **Connection Events**.

Step 8 In the **Remediation Exemption Rule** section, you can view the flows that are exempt from the remediation action.

Step 9 Click **OK** to save the elephant flow settings.

Step 10 Click **Save** to save the policy.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 259.

After configuring your elephant flow settings, monitor your connection events to see if any flows are detected, bypassed, or throttled. You can view this in the **Reason** field of your connection event. The three reasons for elephant flow connections are:

- Elephant Flow
- Elephant Flow Throttled
- Elephant Flow Trusted

Attention Enabling elephant flow detection alone does not cause generation of connection events for elephant flows. If a connection event is already logged for another reason and the flow is also an elephant flow, then the **Reason** field contains this information. However, to ensure that you are logging all elephant flows, you must enable connection logging in the applicable access control rules.

Refer to Cisco Secure Firewall Elephant Flow Detection for more information.

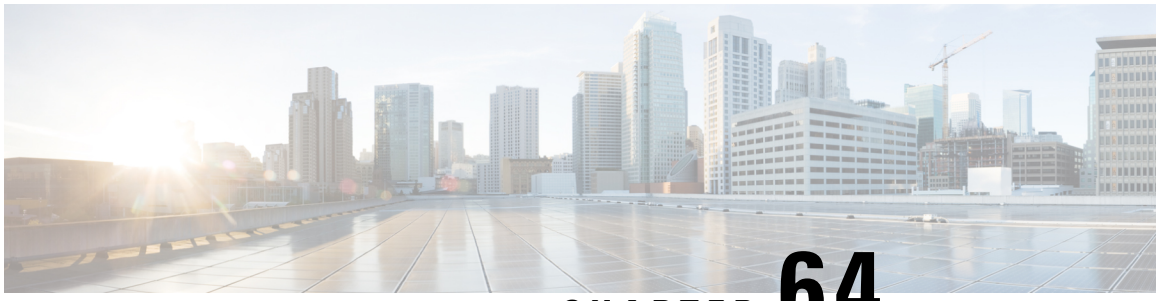

CHAPTER 64

Use Case - Migrate from Snort 2 to Snort 3 In Secure Firewall Management Center

- Migrate from Snort 2 to Snort 3, on page 2159
- Benefits of Migrating to Snort 3, on page 2159
- Sample Business Scenario, on page 2160
- Best Practices for Migrating from Snort 2 to Snort 3, on page 2160
- Prerequisites, on page 2160
- End-to-End Migration Workflow, on page 2160
- Enable Snort 3 on Threat Defense, on page 2161
- Convert Snort 2 Rules of a Single Intrusion Policy to Snort 3, on page 2162
- Deploy Configuration Changes, on page 2167

Migrate from Snort 2 to Snort 3

Snort is an intrusion detection and prevention system that has undergone a significant change from Version 2 to Version 3. To leverage the enhanced features and capabilities of Snort 3, migration of the existing rule sets from Snort 2 becomes crucial. This migration process involves converting and adapting the Snort 2 rules to the Snort 3 rule syntax and optimizing them for improved detection and performance.

In some cases, organizations can have the threat defense devices managed by the Secure Firewall Management Center. Organizations can opt for a hybrid deployment approach during the migration from Snort 2 to Snort 3. This approach allows for a gradual transition and minimizes potential disruptions, if any.

Benefits of Migrating to Snort 3

- **Enhanced protocol support**—Snort 3 provides improved protocol support, allowing you to monitor and detect threats across a wide range of modern protocols, including encrypted traffic.
- **Streamlined rule management**—Snort 3 offers a more user-friendly rule language and rule management system, making it easier to create, modify, and manage rules effectively.
- **Improved performance**—Snort 3 has been optimized to handle higher traffic volumes more efficiently, reducing the risk of performance bottlenecks and ensuring timely threat detection.

Sample Business Scenario

Alice works as a security analyst in a large organization that heavily relies on the Snort inspection engine to monitor and protect their network infrastructure. The organization has been using Snort Version 2 for several years, but they have encountered some limitations and challenges.

Bob, the network administrator, is looking to migrate from Snort 2 to Snort 3 to overcome these issues and enhance his organization's network security capabilities.

This migration will also improve network security monitoring, enhance performance, and streamline rule management.

Best Practices for Migrating from Snort 2 to Snort 3

- Back up your intrusion policy before performing the migration. See the Export Configurations task in the *Cisco Secure Firewall Management Center Administration Guide*.
- Before upgrading a device to Snort 3, if changes are made in Snort 2, use the synchronize utility to include the latest synchronization from Snort 2 to Snort 3 so that you can start with a similar coverage. See Synchronize Snort 2 Rules with Snort 3 , on page 2054.
- Snort 2 custom rules are not automatically converted to Snort 3 and must be manually migrated. See Convert Snort 2 Custom IPS Rules to Snort 3, on page 2052.
- Synchronization does not migrate Snort 2 rules with thresholds or suppressions. These rules must be created again in Snort 3.

Prerequisites

- Have a working knowledge of Snort. To learn about the Snort 3 architecture, see Snort 3 Adoption.
- Back up your management center. See Backup the Management Center.
- Back up your intrusion policy. See Exporting Configurations.

End-to-End Migration Workflow

The following flowchart illustrates the workflow for migrating Snort 2 to Snort 3 in Secure Firewall Management Center.

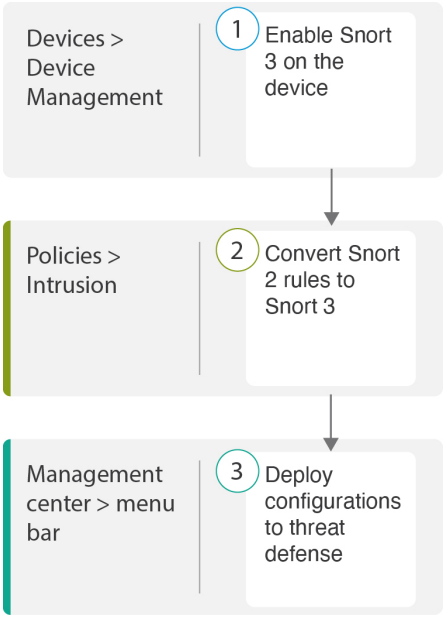

Step	Description
1	Enable Snort 3 on the device. See Enable Snort 3 on Threat Defense, on page 2161.
2	Convert Snort 2 rules to Snort 3. See Convert Snort 2 Rules of a Single Intrusion Policy to Snort 3, on page 2162.
3	Deploy configuration. See Deploy Configuration Changes, on page 2055.

Enable Snort 3 on Threat Defense

Attention During the deployment process, there could be a momentary traffic loss because the current inspection engine needs to be shut down.

Procedure

Step 1 Choose **Devices** > **Device Management**.

Step 2 Click the corresponding device to go to the device home page.

Step 3 Click the **Device** tab.

Step 4 In the **Inspection Engine** section, click **Upgrade**.

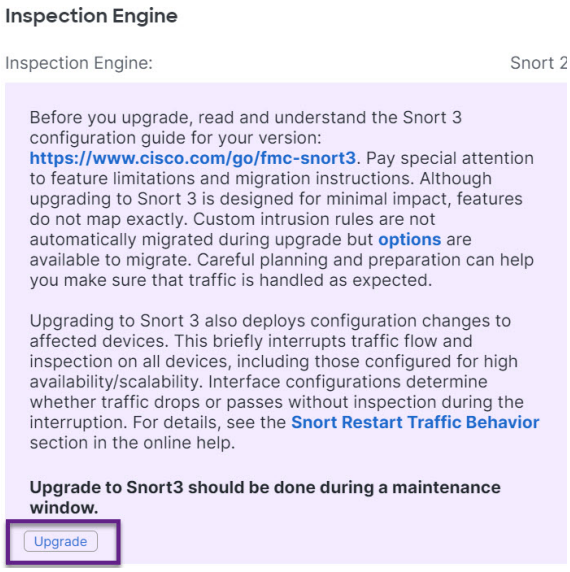

Step 5 Click **Yes**.

What to do next

Deploy the changes on the device. See Deploy Configuration Changes, on page 2055.

The system converts your policy configurations during the deployment process to make them compatible with the selected Snort version.

Convert Snort 2 Rules of a Single Intrusion Policy to Snort 3

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 In the **Intrusion Policies** tab, click **Show Snort 3 Sync status**.

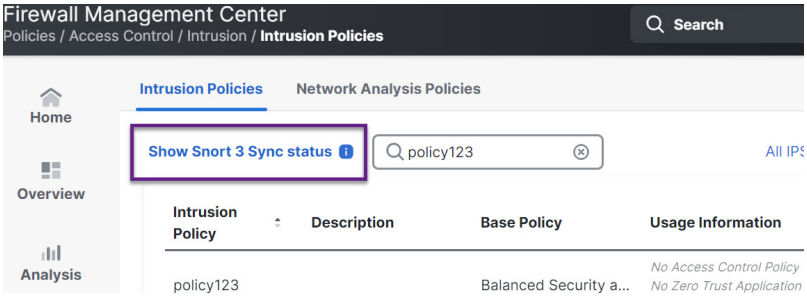

If your policy displays an orange arrow, it indicates that the Snort 2 and the Snort 3 versions of the intrusion policy are not synchronized.

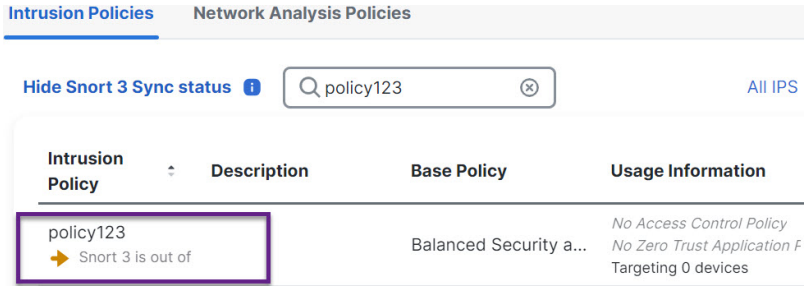

Step 3 Click the orange arrow.

The **Snort 2 to Snort 3 Sync Summary** page displays that the Snort 2 to Snort 3 sync is pending.

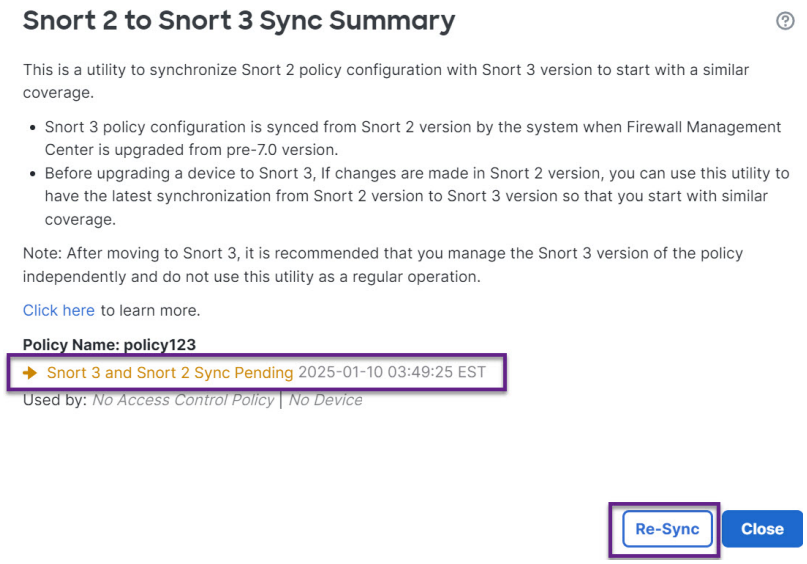

Step 4 Click **Re-Sync** to start the synchronization.

Note
When you click **Re-Sync**, the snort2Lua tool converts the rules from Snort 2 to Snort 3.

The **Summary Details** section lists the rules that were migrated or skipped. In our use case, there are 76 custom Snort 2 rules, 17 rules with thresholds, and 15 rules with suppression that were skipped during the sync process. To migrate the custom rules, go to the next step.

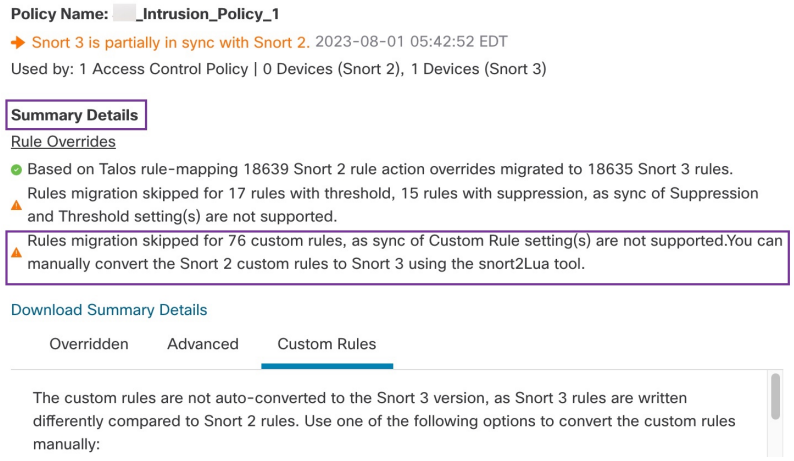

To migrate rules with thresholds and suppressions, go to Step 6.

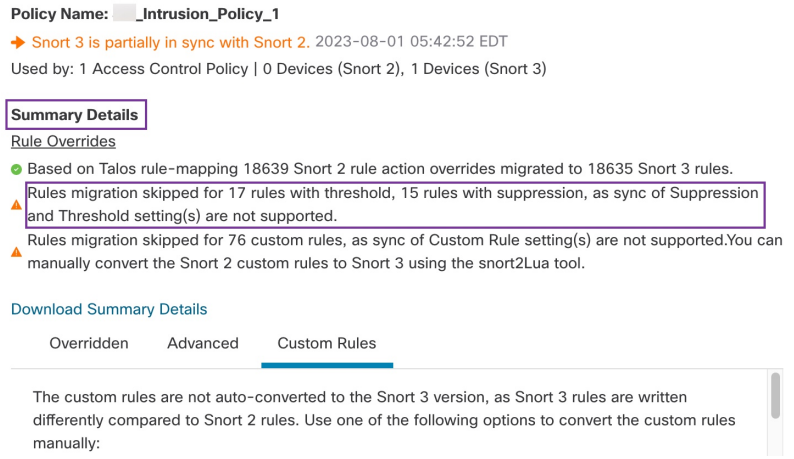

Step 5 To migrate the 76 custom rules, perform either one of these steps:

- In the **Custom Rules** tab, click the **Import** icon to convert and auto-import the local rules to the Snort 3 version of the policy.

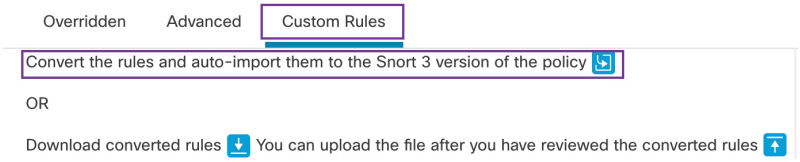

  A confirmation message is displayed after the rules are successfully imported.

- Choose **Objects** > **Intrusion Rules** and click **Snort 3 All Rules**.
  a. Click **Local Rules** in the left panel to check if any rules have been migrated. Notice that no custom rules from Snort 2 have been migrated.
  b. From the **Tasks** drop-down list, choose **Convert Snort 2 rules and import**.

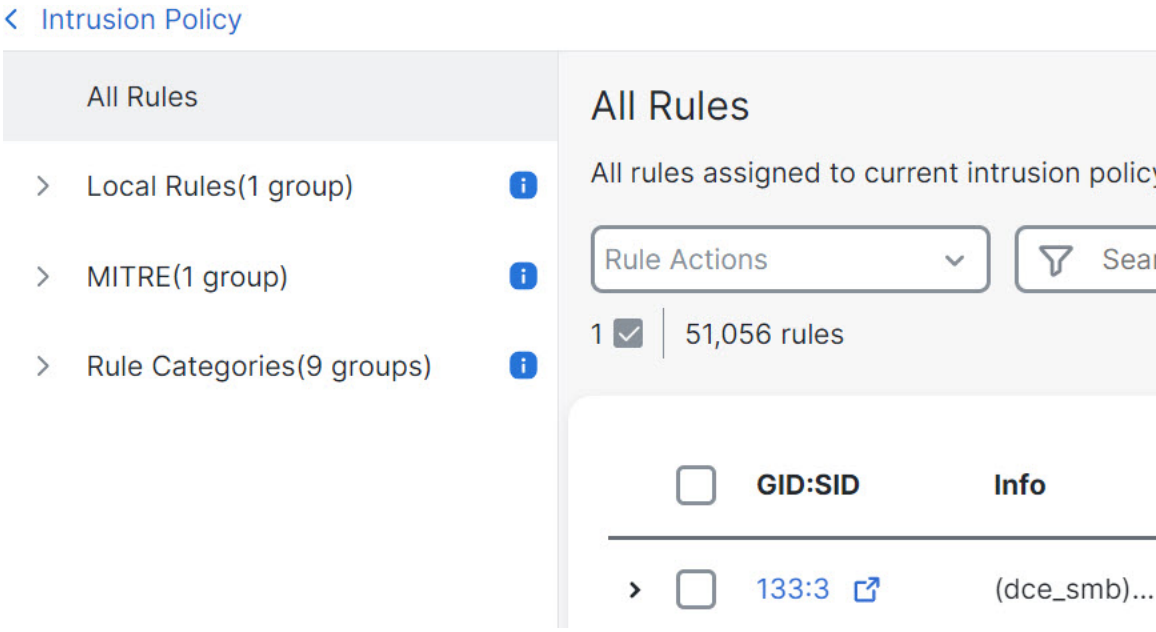

c. Click **OK**.

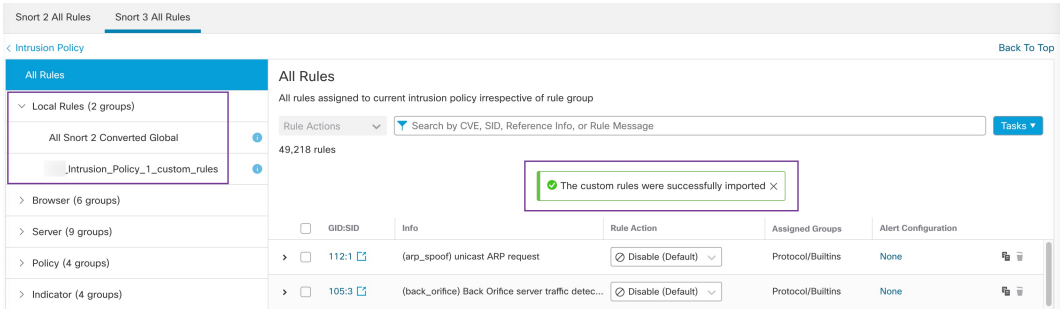

A newly created rule group (**All Snort 2 Converted Global**) is created under **Local Rules** in the left panel.

Notice that all 76 custom rules have been migrated, as shown in the following figure.

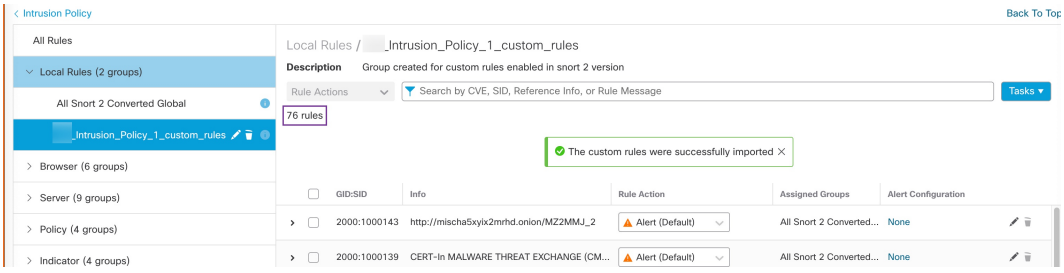

Alternatively, you can select the **Convert Snort 2 rules and download** in the previous step to save the rules file locally. You can review the converted rules in the downloaded file and later upload them using the **Upload Snort 3 rules** option.

Step 6 Click the **Download Summary Details** link to download the rules in .txt format.

The following is a sample of the summary that is displayed.


```
  "id": "00505691-15DC-0ed3-0000-004294988561",
  "name": "_Intrusion_Policy_1",
  "type": "IntrusionPolicy",
  "syncStatus": {
    "source": {
      "id": "bdce2d6a-1ebe-11ee-8e88-220032eb1fb5",
      "type": "IntrusionPolicy"
    },
    "status": "WARN",
    "description": "Migration is partially successful. Some of the rules are not copied to
 Snort3.",
    "timestamp": 1690883954814,
    "lastUser": {
      "name": "admin"
    },
    "details": [
      {
        "type": "Summary",
        "status": "INFO",
        "description": "Based on Talos rule-mapping 18639 Snort 2 rule action overrides
migrated to 18635 Snort 3 rules."
      },
      {
        "id":
"1:1000156=alert,1:1000114=alert,1:1000160=alert,1:1000135=alert,1:1000115=alert,1:1000118=alert,

1:1000092=alert,1:1000139=alert,1:1000123=alert,1:1000159=alert,1:1000149=disabled,1:1000167=alert,

1:1000133=alert,1:1000095=alert,1:1000143=alert,1:1000106=alert,1:1000153=alert,1:1000097=alert,1:1000141=alert,

1:1000148=alert,1:1000090=alert,1:1000119=alert,1:1000112=alert,1:1000138=alert,1:1000128=alert,1:1000132=alert,

1:1000134=alert,1:1000145=disabled,1:1000110=disabled,1:1000107=alert,1:1000163=alert,1:1000124=alert,1:1000125=alert,

1:1000094=alert,1:1000113=disabled,1:1000147=alert,1:1000161=alert,1:1000105=disabled,1:1000140=alert,1:1000111=alert,

1:1000102=alert,1:1000129=disabled,1:1000108=alert,1:1000144=disabled,1:1000088=alert,1:1000091=alert,1:1000131=alert,

1:1000157=alert,1:1000120=alert,1:1000126=alert,1:1000165=alert,1:1000146=alert,1:1000162=alert,1:1000116=alert,1:1000142=alert,

1:1000170=disabled,1:1000169=alert,1:1000104=alert,1:1000099=disabled,1:1000171=alert,1:1000093=alert,1:1000087=alert,1:1000100=alert,

1:1000137=alert,1:1000158=alert,1:1000103=alert,1:1000098=alert,1:1000127=disabled,1:1000130=alert,1:1000164=alert,1:1000089=alert,

         1:1000109=alert,1:1000136=alert,1:1000117=alert,1:1000166=alert,1:1000168=alert",
        "type": "PolicyInfo",
        "description": "Corresponding Snort 2 policy overridden custom (local) rules."
      },
      {
        "type": "AssignedDevices",
        "status": "INFO",
        "description": "Snort3:0 , Snort2:0"
      },
```

```
{
  "id": "122:6",
  "type": "Threshold",
  "status": "ERROR",
  "description": "PSNG_TCP_FILTERED_DECOY_PORTSCAN"
},
{
  "id": "122:15",
  "type": "Threshold",
  "status": "ERROR",
  "description": "PSNG_IP_PORTSWEEP_FILTERED"
 },
{
  "id": "122:1",
  "type": "Threshold",
  "status": "ERROR",
  "description": "PSNG_TCP_PORTSCAN"
},
```


Step 7 Click **Close** to close the **Sync Summary** dialog box.

Step 8 To check the rules with status: ERROR, choose **Policies** > **Intrusion** and click the **Snort 2** version of the intrusion policy.

Step 9 Under **Policy Information**, click **Rules** and filter for the rule. For example, enter **PSNG_TCP_PORTSCAN** in the **Filter**field to find the rule.

Step 10 Click **Show Details** to view the detailed version of the rule.

Step 11 Create the rule again in Snort 3 using Snort 3 rule guidelines and save the file as a .txt or .rules file. For more information, see www.snort3.org.

Step 12 Upload the custom rule that you just created locally to the list of all the Snort 3 rules. See Add Custom Rules to Rule Groups.

What to do next

Deploy configuration changes. See Deploy Configuration Changes, on page 2055.

Deploy Configuration Changes

After you change configurations, deploy them to the affected devices.

Note This topic covers the basic steps involved in deploying configuration changes. We *strongly* recommend that you refer the *Deploy Configuration Changes* topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide* to understand the prerequisites and implications of deploying the changes before proceeding with the steps.

Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic.

Procedure

Step 1 On the Secure Firewall Management Center menu bar, click **Deploy** and choose **Deployment**.

The GUI page lists the devices with out-of-date configurations having **Pending**status.

- The **Modified By** column lists the users who have modified the policies or objects. Expand the device listing to view the users who have modified the policies for each policy listing.

  Note
  Usernames are not provided for deleted policies and objects.

- The **Inspect Interruption** column indicates if traffic inspection interruption might occur in the device during deployment.

  If this column is blank for a device, it indicates that there will be no traffic inspection interruptions on that device during deployment.

- The **Last Modified Time** column specifies the last time you made configuration changes.
- The **Preview** column allows you to preview the changes for the next deployment.
- The **Status** column provides the status for each deployment.

Step 2 Identify and choose the devices on which you want to deploy configuration changes.

- Search—Search for the device name, type, domain, group, or status in the search box.
- Expand—Click **Expand Arrow** (>) to view device-specific configuration changes to be deployed.

  When you check a check box adjacent to a device, all the changes made to the device and listed under the device, are pushed for deployment. However, you can use **Policy selection** ( ) to select individual policies or specific configurations to deploy while withholding the remaining changes without deploying them.

  Note
  - When the status in the **Inspect Interruption** column indicates (**Yes**) that deploying will interrupt inspection, and perhaps traffic, on a Firewall Threat Defense device, the expanded list indicates the specific configurations causing the interruption with the **Inspect Interruption** ( ).
  - When there are changes to interface groups, security zones, or objects, the impacted devices are shown as out-of-date on the Firewall Management Center. To ensure that these changes take effect, the policies with these interface groups, security zones, or objects, also need to be deployed along with these changes. The impacted policies are shown as out-of-date on the **Preview**page on the Firewall Management Center.

Step 3 Click **Deploy**.

Step 4 If the system identifies errors or warnings in the changes to be deployed, it displays them in the **Validation Messages** window. To view complete details, click the arrow icon before the warnings or errors.

You have the following choices:

- Deploy—Continue deploying without resolving warning conditions. You cannot proceed if the system identifies errors.

- Close—Exit without deploying. Resolve the error and warning conditions, and attempt to deploy the configuration again.

What to do next

During deployment, if there is a deployment failure, there is a possibility that the failure may impact traffic. However, it depends on certain conditions. If there are specific configuration changes in the deployment, the deployment failure may lead to traffic being interrupted. For details, see the Deploy Configuration Changes topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide*.

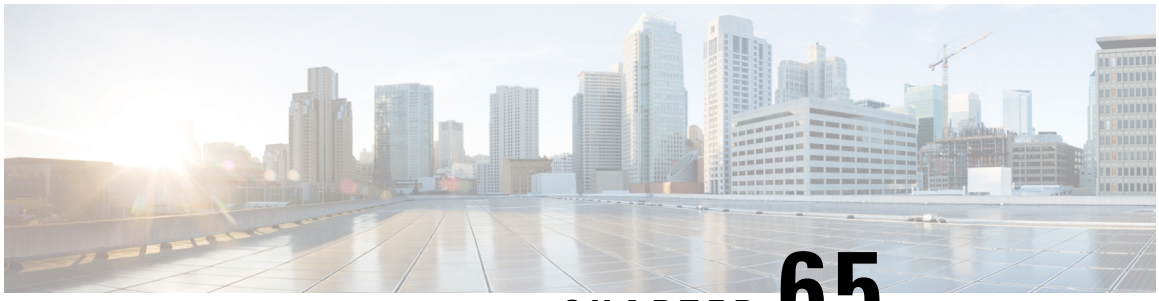

CHAPTER 65

Use Case - Generate Snort 3 Recommendations In Secure Firewall Management Center

- Snort 3 Rule Recommendations, on page 2171
- Benefits, on page 2172
- Sample Business Scenario, on page 2172
- Best Practices, on page 2172
- Prerequisites, on page 2172
- Generate Snort 3 Recommendations, on page 2172
- Deploy Configuration Changes, on page 2175

Snort 3 Rule Recommendations

Rule recommendations automatically tune your intrusion policy with rules that are specific to the host environment. You can enable additional rules or tune the current rule set by disabling rules for the vulnerabilities that are not present in your network. For more information, see Overview of Secure Firewall Recommended Rules, on page 2093.

How does it work?

The management center builds a database of hosts on your network with details such as the IP address, hostname, operating system, services, users, and client applications through passive discovery. Based on this information, the system maps vulnerabilities to each discovered host. The Recommendations feature uses this host database to determine the rules that apply to your environment.

In Snort 3, there are four security levels, each corresponding to a specific Talos policy. They are:

- Level 1–Connectivity Over Security
- Level 2–Balanced Security and Connectivity
- Level 3–Security Over Connectivity
- Level 4–Maximum Detection

Check the **Accept Recommendations to Disable Rules** check box to disable rules for vulnerabilities not found on the hosts in your network. Check this option only if you have to trim your rule set because of a high number of alerts, or to improve inspection performance.

Benefits

- By configuring recommendations, you can tailor your intrusion policy to detect specific types of threats more effectively using rules that are specific to the host environment.
- Recommendations contribute to a more efficient and effective incident response process by reducing false positives and false negatives.

Sample Business Scenario

A large corporate network uses Snort 3 as its primary intrusion detection and prevention system. In a rapidly evolving threat landscape, robust network security measures must be adopted. The security team wants to enhance their incident response capabilities. One of the ways to do that is to generate recommendations or rule sets based on the vulnerabilites detected in the host network. This helps to optimize their intrusion policies, thereby safeguarding the network more effectively.

Best Practices

- You must have quality accurate host data.

  Because of the passive nature of Network Discovery, your threat defense devices must be positioned as close as possible to your protected hosts. This allows the threat defense devices to watch network traffic to and from these hosts, giving you an accurate data about applications, services, and vulnerabilities present on your network.
- Devices should have visibility to East-West as well as North-South traffic flows to build an accurate host profile.
- You can create a scheduled task to update recommendations automatically.

Prerequisites

- Ensure that hosts are present in the system to generate recommendations.
- Protected networks configured for recommendations should map to the hosts present in the system.

Generate Snort 3 Recommendations

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Click the **Snort 3 Version** button of the corresponding intrusion policy.

Step 3 Click the **Recommendations (Not in Use)** layer to configure the rule recommendations.

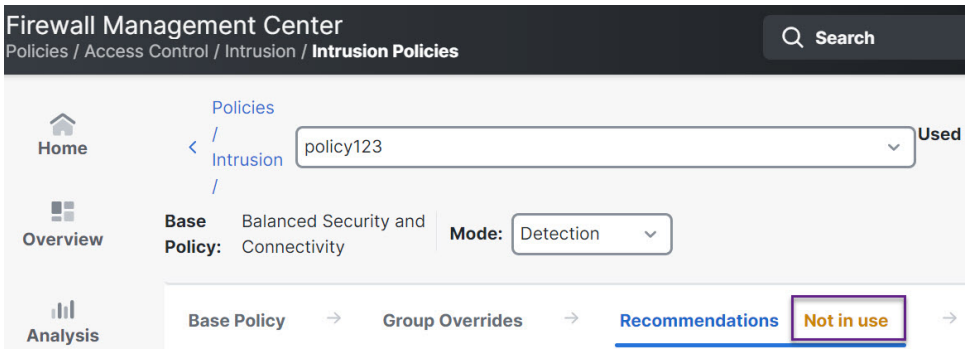

In the **Cisco Recommended Rules** window, you can set the security level.

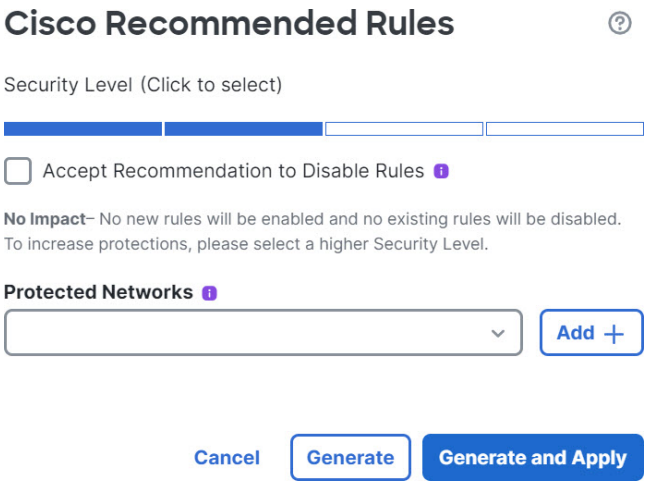

Step 4 Click to select the security level.

Step 5 (Optional) Check the **Accept Recommendation to Disable Rules** check box to disable the rules written for vulnerabilities not found on the hosts in your network.

Use this option, only if you have to trim your rule set because of a high number of alerts or to improve inspection performance.

Step 6 From the **Protected Networks** drop-down list, choose the network objects that must be examined by the recommendations. By default, any IPv4 or IPv6 networks are selected if you do not make a selection.

Click **Add +** to create a new network object of type **Host** or **Network** and click **Save**.

Step 7 Generate and apply recommendations:

- **Generate**—Generates the recommendations for an intrusion policy. This action lists the rules under **Recommended Rules (Not in use)**.
- **Generate and Apply**—Generates and applies the recommendations for an intrusion policy. This action lists the rules under **Recommended Rules (Not in use)**.

Recommendations are generated successfully. A new recommendation tab appears with all the recommended rules and their corresponding recommended actions. Rule action preset filters are also available for this tab, in addition to new recommendations.

Step 8 Verify the recommendations and then apply them accordingly:

- **Accept**—Applies the previously generated recommendations for an intrusion policy.
- **Refresh**—Regenerates and updates the rule recommendations for an intrusion policy.
- **Edit**—Opens the **Recommendations** dialog box where you can provide the recommendation input values and then generate the recommendations.
- **Discard**—Either reverts or removes the applied recommended rules from the policy; also removes the **Recommendations** tab.

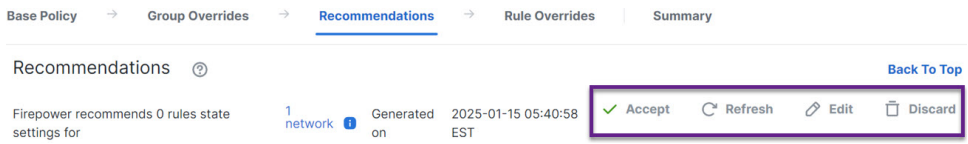

Under **All Rules**, the Recommended Rules section displays the recommended rules.

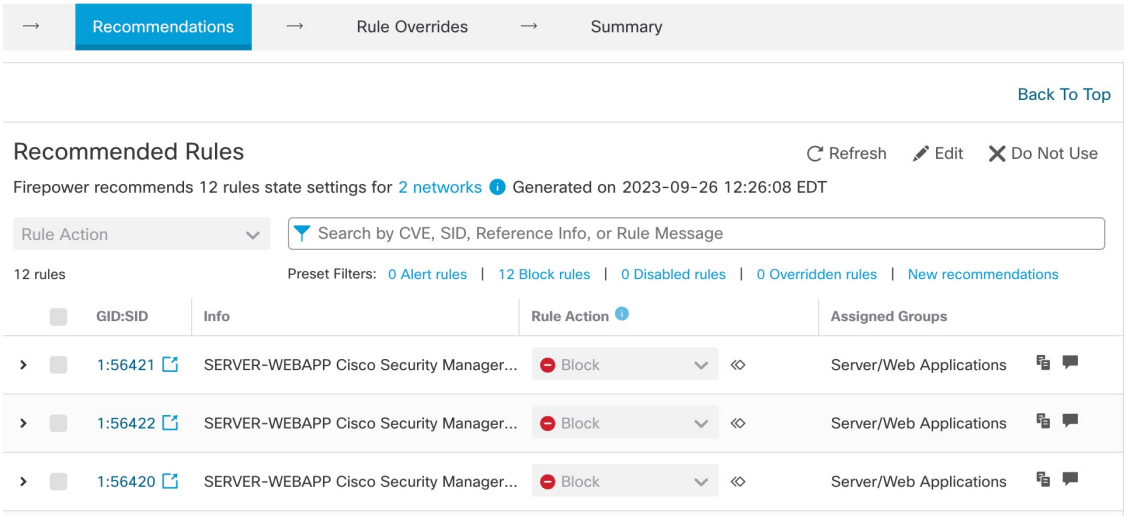

Step 9 To effectively use recommendations, they must be updated periodically. Follow these steps:

a. Choose **System** (⚙) > **Tools** > **Scheduling**.

b. Click **Add Task**.

c. Choose **Cisco Recommended Rules** from the **Job Type** drop-down list.

d. Update the required fields, as needed.

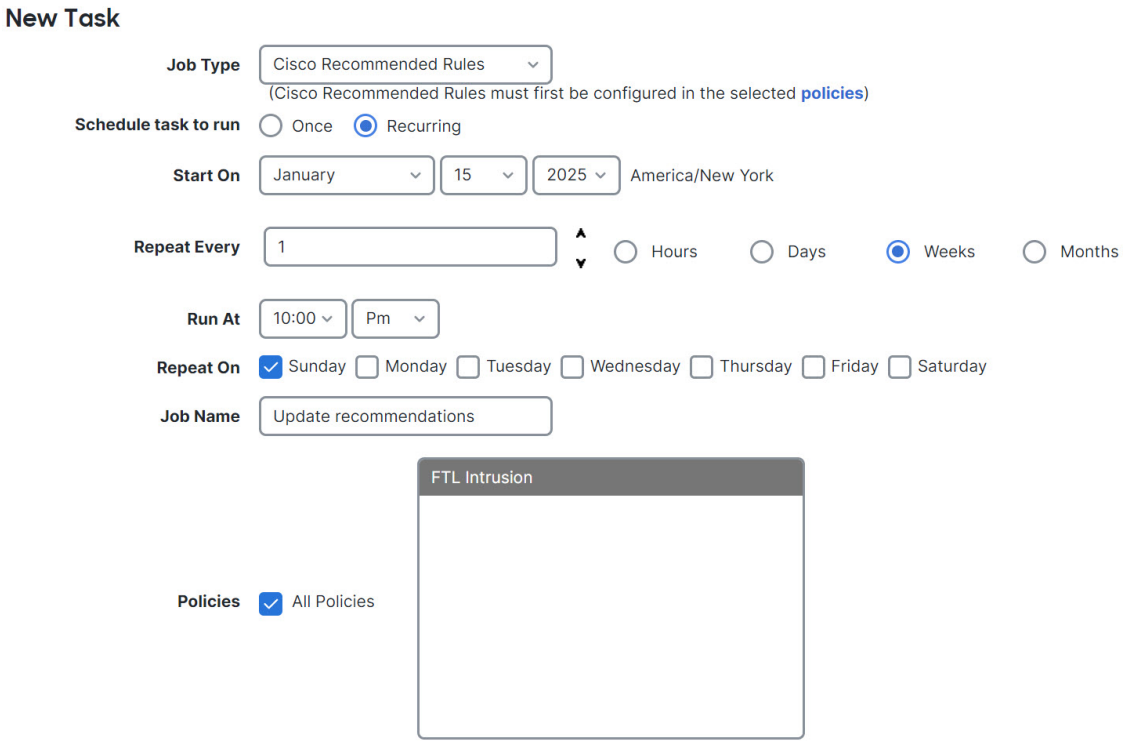

e. Click **Save**.

What to do next

Deploy configuration changes. See Deploy Configuration Changes, on page 2055.

Deploy Configuration Changes

After you change configurations, deploy them to the affected devices.

Note This topic covers the basic steps involved in deploying configuration changes. We *strongly* recommend that you refer the *Deploy Configuration Changes* topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide* to understand the prerequisites and implications of deploying the changes before proceeding with the steps.

Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic.

Procedure

Step 1 On the Secure Firewall Management Center menu bar, click **Deploy** and choose **Deployment**.

The GUI page lists the devices with out-of-date configurations having **Pending**status.

- The **Modified By** column lists the users who have modified the policies or objects. Expand the device listing to view the users who have modified the policies for each policy listing.

  Note
  Usernames are not provided for deleted policies and objects.

- The **Inspect Interruption** column indicates if traffic inspection interruption might occur in the device during deployment.

  If this column is blank for a device, it indicates that there will be no traffic inspection interruptions on that device during deployment.

- The **Last Modified Time** column specifies the last time you made configuration changes.
- The **Preview** column allows you to preview the changes for the next deployment.
- The **Status** column provides the status for each deployment.

Step 2 Identify and choose the devices on which you want to deploy configuration changes.

- Search—Search for the device name, type, domain, group, or status in the search box.
- Expand—Click **Expand Arrow** (>) to view device-specific configuration changes to be deployed.

  When you check a check box adjacent to a device, all the changes made to the device and listed under the device, are pushed for deployment. However, you can use **Policy selection** ( ) to select individual policies or specific configurations to deploy while withholding the remaining changes without deploying them.

  Note
  - When the status in the **Inspect Interruption** column indicates (**Yes**) that deploying will interrupt inspection, and perhaps traffic, on a Firewall Threat Defense device, the expanded list indicates the specific configurations causing the interruption with the **Inspect Interruption** ( ).
  - When there are changes to interface groups, security zones, or objects, the impacted devices are shown as out-of-date on the Firewall Management Center. To ensure that these changes take effect, the policies with these interface groups, security zones, or objects, also need to be deployed along with these changes. The impacted policies are shown as out-of-date on the **Preview**page on the Firewall Management Center.

Step 3 Click **Deploy**.

Step 4 If the system identifies errors or warnings in the changes to be deployed, it displays them in the **Validation Messages** window. To view complete details, click the arrow icon before the warnings or errors.

You have the following choices:

- Deploy—Continue deploying without resolving warning conditions. You cannot proceed if the system identifies errors.

- Close—Exit without deploying. Resolve the error and warning conditions, and attempt to deploy the configuration again.

What to do next

During deployment, if there is a deployment failure, there is a possibility that the failure may impact traffic. However, it depends on certain conditions. If there are specific configuration changes in the deployment, the deployment failure may lead to traffic being interrupted. For details, see the Deploy Configuration Changes topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide*.

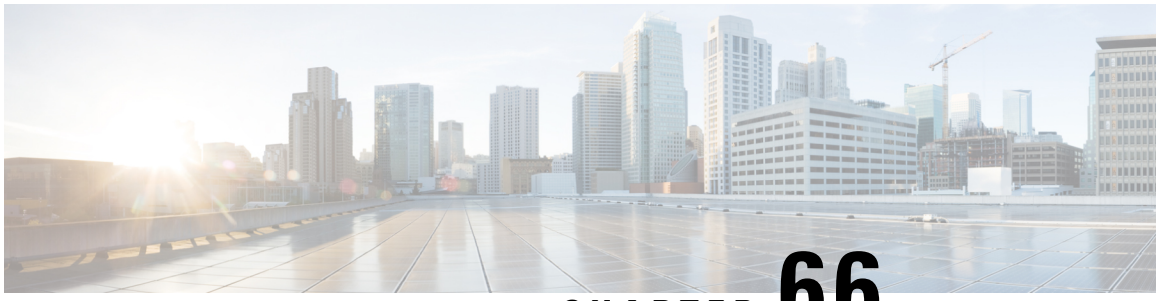

CHAPTER 66

Use Case - Block Traffic Based on the EVE Threat Confidence Score

- About Encrypted Visibility Engine, on page 2179
- Benefits, on page 2179
- Sample Business Scenario, on page 2179
- Prerequisites, on page 2180
- High-Level Workflow, on page 2180
- Configure Block Thresholds in EVE, on page 2180
- Additional References, on page 2185

About Encrypted Visibility Engine

You can use the Encrypted Visibility Engine (EVE) to identify client applications and processes using Transport Layer Security (TLS) encryption. EVE provides more visibility into the encrypted sessions without decryption. Based on EVE's findings, administrators can enforce policy actions on the traffic within their environments. You can also use the EVE to identify and stop malware.

Benefits

Administrators can leverage and adjust EVE's threat score to block malicious encrypted traffic. If the probability that the incoming traffic is malicious, then based on the threat score, you can configure EVE to block the connection.

Sample Business Scenario

A large corporate network uses Snort 3 as its primary intrusion detection and prevention system. In a rapidly evolving threat landscape, adoption of robust network security measures is necessary and important. The security team uses EVE to enhance encrypted traffic inspection without the need to implement full man-in-the-middle (MITM) decryption. The EVE technology uses fingerprints of known malicious processes to identify and stop malware. Network administrators must have the flexibility to configure EVE's block traffic thresholds to block potentially malicious connections, which are based on their configured block thresholds.

Prerequisites

- You must be running management center 7.4.0 or later, and the managed threat defense must also be 7.4.0 or later.
- Ensure that you have a valid Intrusion Prevention System (IPS) license and Snort 3 is the detection engine.

High-Level Workflow

1. EVE analyzes the incoming traffic and gives a verdict on the probability of incoming traffic being malware or not.
2. If EVE detects incoming traffic to be malware with a certain level of confidence, you can configure EVE to block that traffic.
3. The packets are first checked for malware probability or threat score, and the threat score is compared with the block threshold that you have set.
4. If the threat score is higher than the configured threshold, EVE blocks the traffic.
5. If the threat score is lesser than the configured threshold, EVE takes no action.

Configure Block Thresholds in EVE

This procedure shows how to block potentially malicious traffic, based on the EVE threat confidence score of 90 percent or higher.

Procedure

Step 1 Choose **Policies** > **Access Control**.

Step 2 Click **Edit** (✎) next to the access control policy you want to edit.

Step 3 Choose **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 4 Click **Edit** (✎) next to **Encrypted Visibility Engine**.

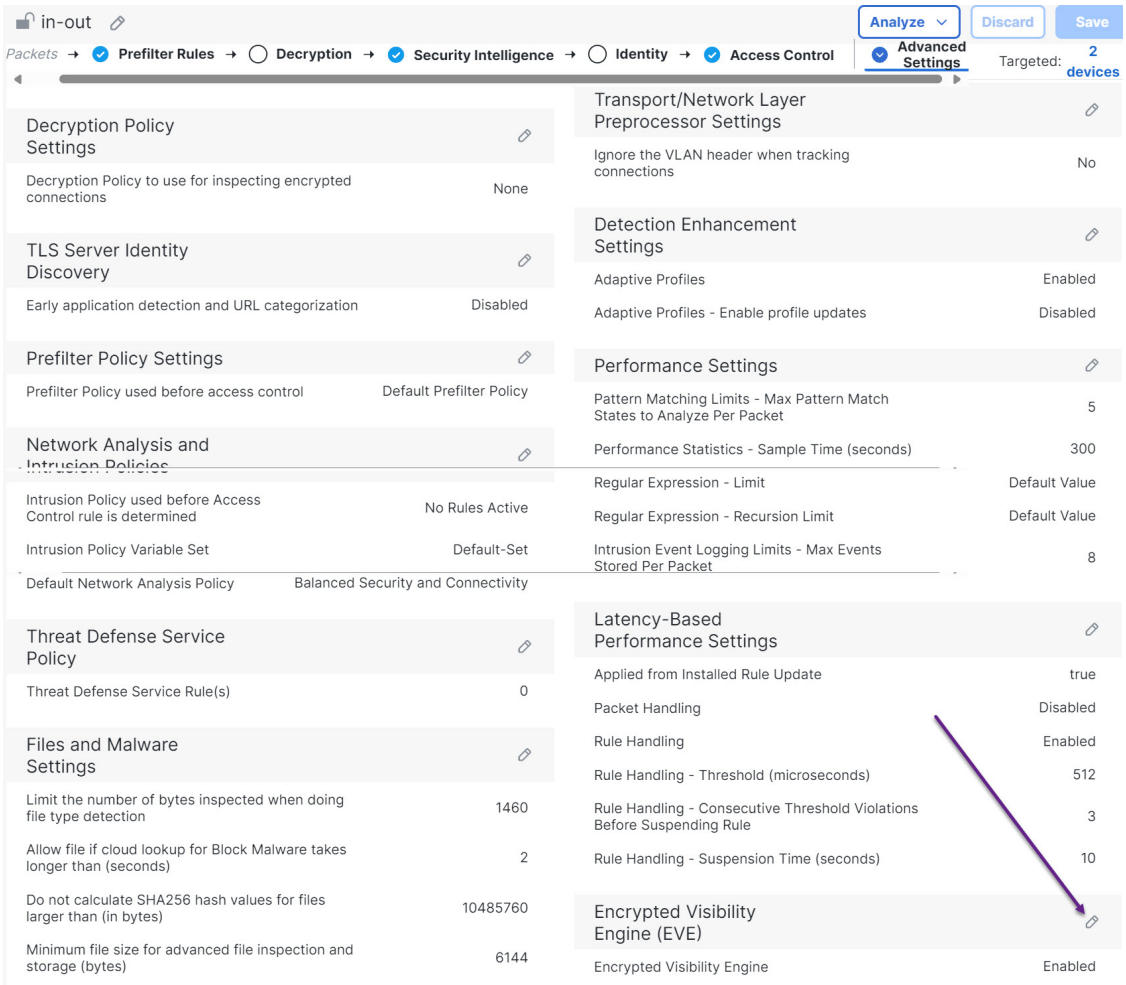

Step 5 In the **Encrypted Visibility Engine** page, enable the **Encrypted Visibility Engine (EVE)** toggle button.

Step 6 Enable the **Block Traffic Based on EVE Score** toggle button. Any incoming traffic that is a potential threat is blocked by default.

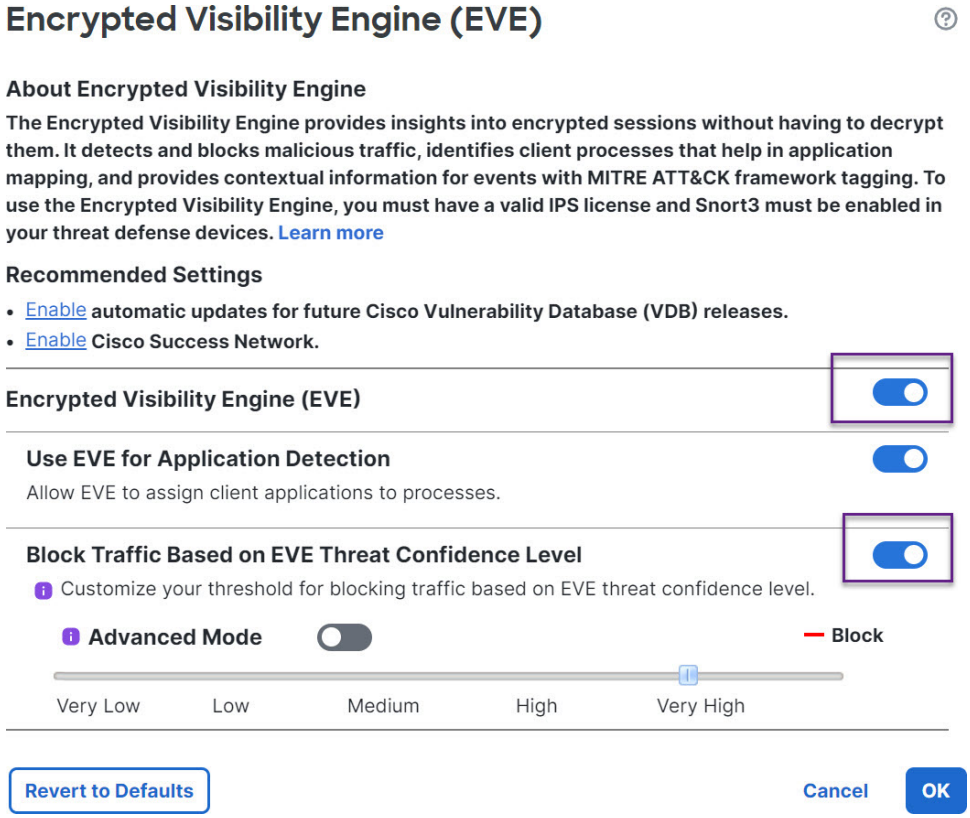

Note

By default, the threshold at which malware is blocked is 99 percent, which means:

- If EVE detects the traffic to be malware with 99 percent confidence or higher, EVE blocks the traffic.
- If EVE detects the traffic to be malware with less than 99 percent confidence, EVE takes no action.

Step 7 Use the slider to adjust the threshold for blocking based on EVE threat confidence. This ranges from **Very Low** to **Very High**. In this example, the slider is set to **Very High**.

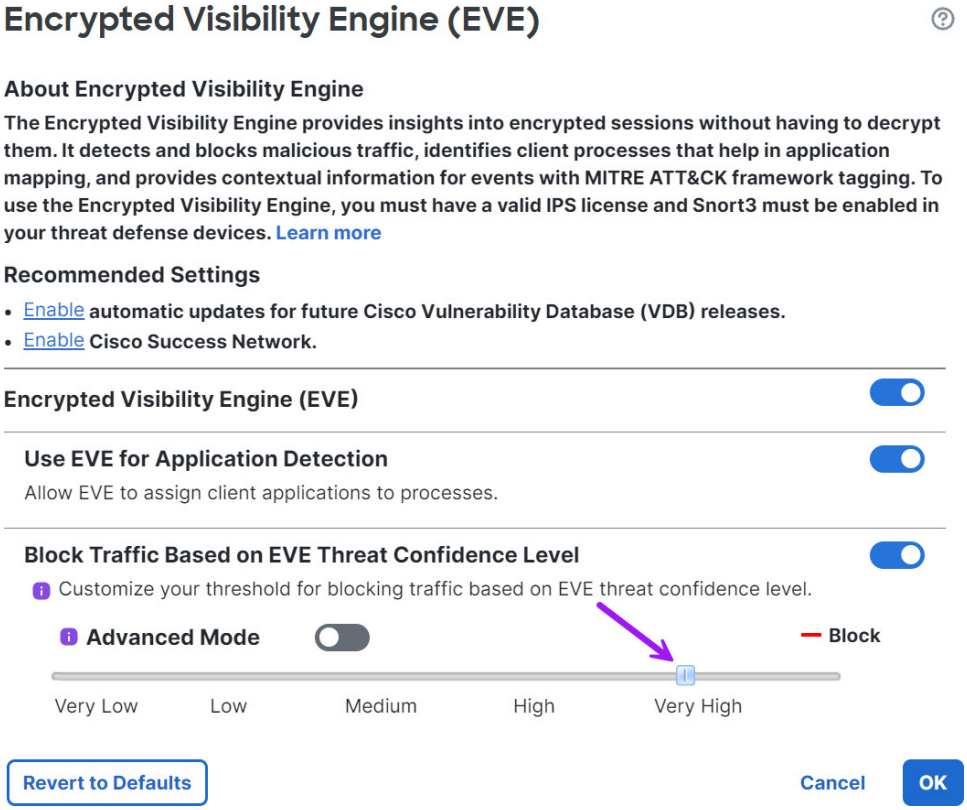

Step 8 For further granular control, enable the **Advanced Mode** toggle button. Now, you can assign a specific EVE Threat Confidence Score for blocking traffic. The default threshold is 99 percent.

Step 9 In this example, change the block threshold to **90** percent.

Attention
As a best practice, we recommend that you do not set the block threshold to below 50 percent to ensure optimum performance.

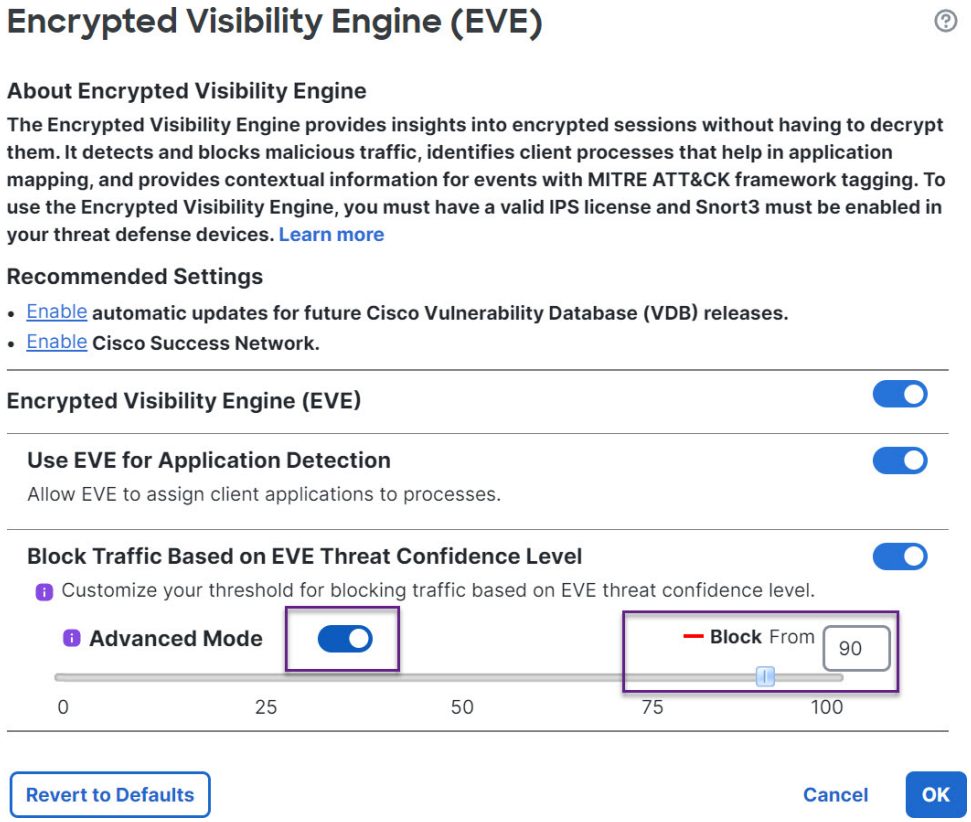

Step 10 Click **OK**.

Step 11 Click **Save**.

What to do next

Deploy configuration changes. See Deploy Configuration Changes, on page 2055.

View EVE Events

Procedure

Step 1 To verify the block action, choose **Analysis** > **Connections**. You can also view the events from the **Unified Events** viewer.

Step 2 If you have configured EVE to block traffic, the **Reason** field shows **Encrypted Visibility Block**.

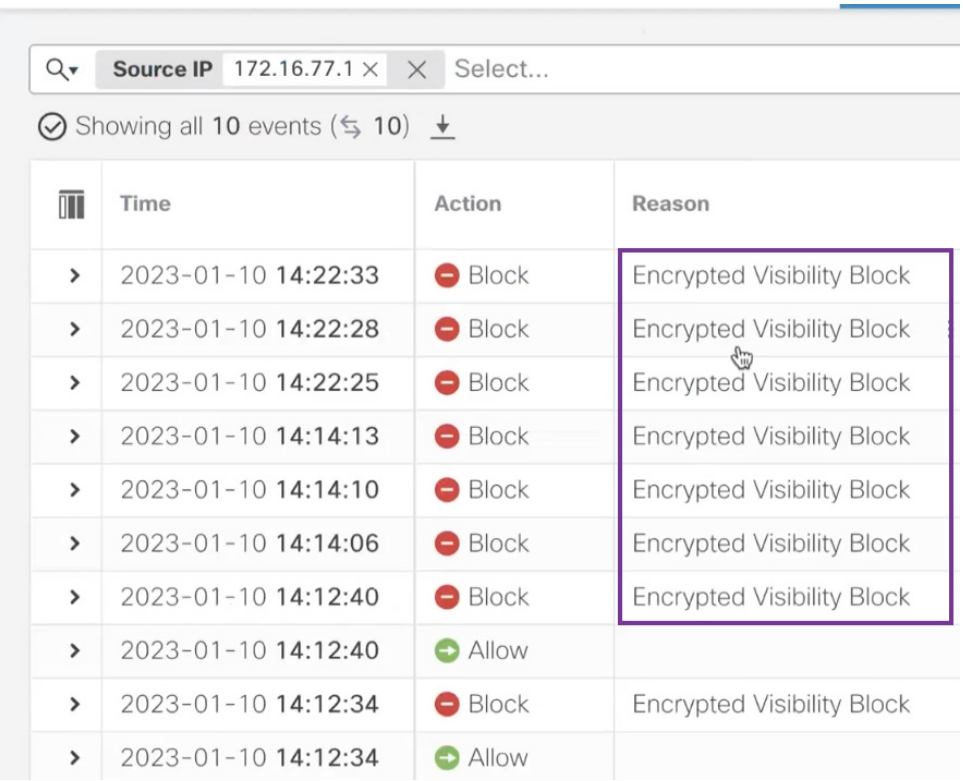

Step 3 The following is an example of the **Encrypted Visibility Process Name** as **test_malware**, **Encrypted Visibility Threat Confidence** as **Very High**, and **Encrypted Visibility Threat Confidence Score** as **90** percent.

Additional References

For detailed conceptual information, see the Encrypted Visibility Engine for Snort 3 chapter in this guide or the content in the following link:

Encrypted Visibility Engine

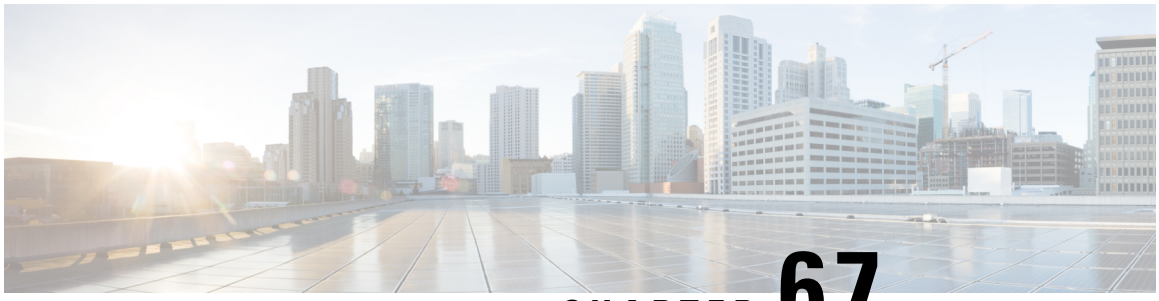

CHAPTER 67

Use Case - Configure Elephant Flow Detection Outcomes

- About Elephant Flows, on page 2187
- Benefits of Elephant Flow Detection and Remediation, on page 2187
- Elephant Flow Workflow, on page 2188
- Sample Business Scenario, on page 2188
- Prerequisites, on page 2189
- Configure Elephant Flow Parameters, on page 2189
- Configure Elephant Flow Remediation Exemption, on page 2193
- Additional References, on page 2196

About Elephant Flows

Elephant flows are extremely large (in total bytes), relative long-running network connections set up by a TCP (or other protocols) flow measured over a network link. By default, elephant flows are flows or connections that are larger than 1 GB per 10 seconds. They can cause performance duress or issues in Snort cores. Elephant flows are important because they can potentially consume an excessive amount of CPU resources and impact other competing flows for detection resources and cause issues, such as increased latency or packet drops.

Benefits of Elephant Flow Detection and Remediation

- Elephant flow configuration allows customization and the option to bypass or even throttle elephant flows.
- You can choose to bypass or throttle flows that are based on your chosen applications to provide Snort inspection of suspect traffic, while bypassing more trusted traffic.
- Elephant flow remediation helps prioritize and free up more bandwidth for your internal applications, depending on your specific requirements.

Elephant Flow Workflow

When an elephant flow is detected based on your configured parameters, you can choose to bypass or throttle the flow. When a flow is bypassed, the traffic is allowed to pass without Snort inspection. Throttling indicates that the flow throughput is reduced. The reduction on flow rate is done in 10 percent increments until the CPU utilization reduces to below the configured threshold. Bypassing or throttling happens after identifying the elephant flow and meeting the additional CPU and time window parameters. Prior to identification of the elephant flow, your intrusion policy processes the flow, assuming that you have configured this in an Allow rule. This means that elephant flows are not allowed to pass through the system completely uninspected because most of the attacks are detected very early in a connection.

To understand how flows are handled, see the following flow diagram.

Figure 449: Elephant Flow Workflow

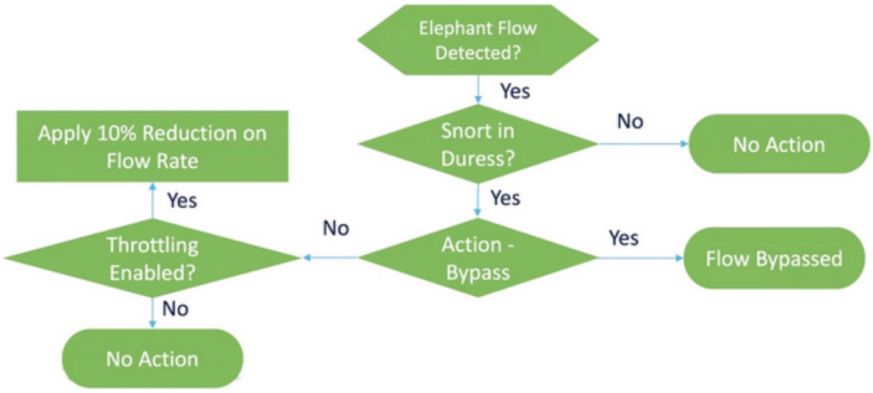

No action is taken unless the system detects a Snort duress condition (performance issue). The system does not throttle or bypass a flow just because it is large. Also, the actions of throttle and bypass are mutually exclusive. This means that you can either bypass or throttle a flow, but not both.

If you do not want to bypass all the elephant flows causing duress, you can limit the bypass option to specific applications only. You can prioritize connectivity for the applications that you trust, without throttling performance. You can configure the applications that must be bypassed, but the remaining flows (causing duress) are throttled. This ensures that the other nontrusted application flows still receive full Snort inspection although their bandwidth is reduced.

Sample Business Scenario

In a data center, several activities are happening, such as replication of data between clusters, virtual machine integration, and database backup. Users in an organization could be watching videos on an OTT or downloading them. Bandwidth utilization for such activities might result in elephant flows, slow down the network, and impact the performance of important tasks. As a network administrator (and depending on your specific requirements), you want visibility into such large flows that are causing bandwidth issues and remediate them.

As an example, let us see how you can configure elephant flow parameters to bypass Snort inspection for WebEx traffic (which your organization uses for real-time video conferencing) and throttle the remaining applications or connections, including videos, movies, and so on.

Prerequisites

- Ensure that you are running management center 7.2.0 or later and that the managed threat defense is also 7.2.0 or later.
- Only enabling elephant flow detection does not generate additional connection events. Elephant flow detection adds the Elephant Flow notation to matching connections that are already being logged to the management center. **To log these events, you must enable connection logging in your access control policy.** You can do that for specific rules or add a Monitor rule that logs all connections, including elephant flows.

Configure Elephant Flow Parameters

Procedure

Step 1 Choose **Policies** > **Access Control**.

Step 2 Click **Edit** (✎) next to the access control policy that you want to edit.

Step 3 Choose **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 4 Click **Edit** (✎) next to **Elephant Flow Settings**.

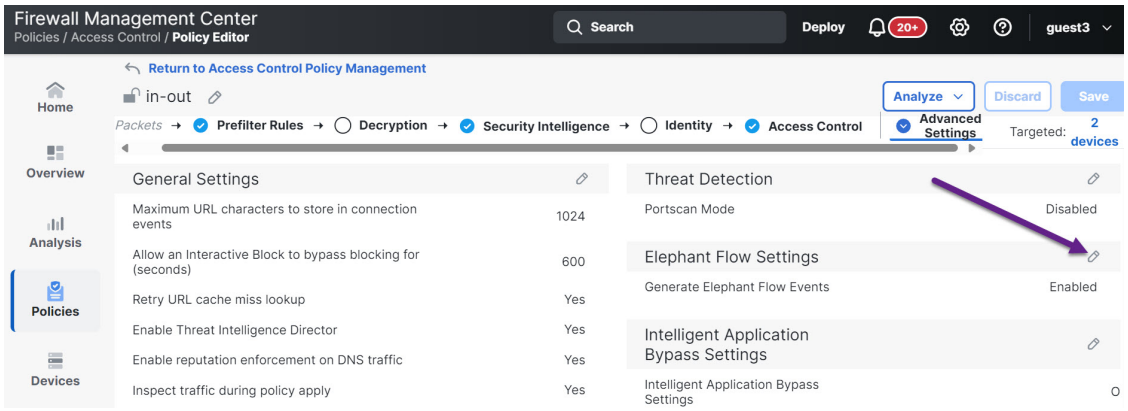

Step 5 The **Elephant Flow Detection** toggle button is enabled by default. The default setting enables detection only and no default action is configured. The detection settings allow you to adjust the flow bytes and duration so that you can identify the elephant flows in your system.

As a test setting, configure the flow bytes and duration parameters, as shown in the following figure.

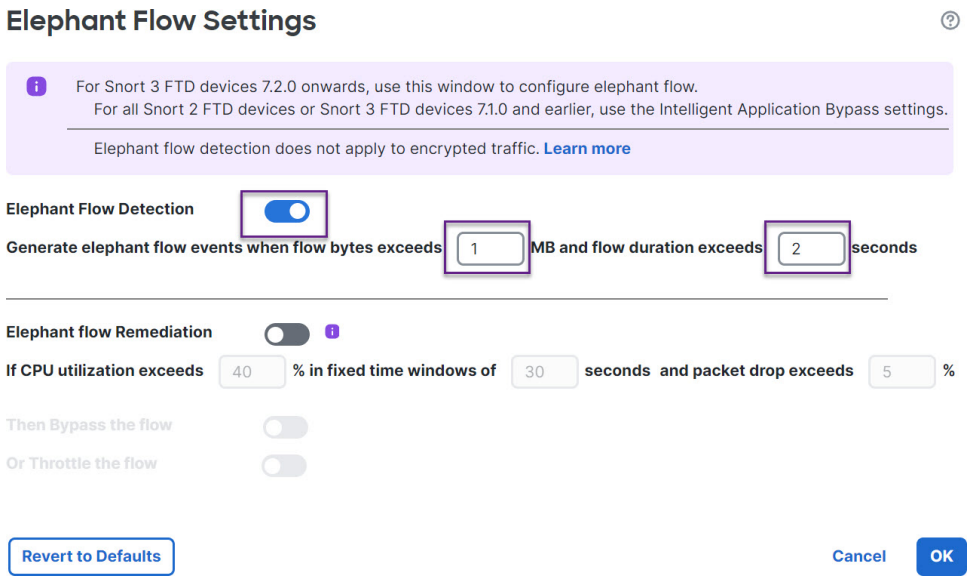

Step 6 Enable the **Elephant Flow Remediation** toggle button. When an elephant flow is detected, you can choose to bypass or throttle the flow. Bypassing a flow means that the traffic is allowed to pass without Snort inspection. Throttling indicates that the flow throughput is reduced. This rate reduction is done in 10 percent increments until the CPU utilization reduces to lesser than the configured threshold.

As a test setting, configure the elephant flow remediation parameters as shown in the following figure.

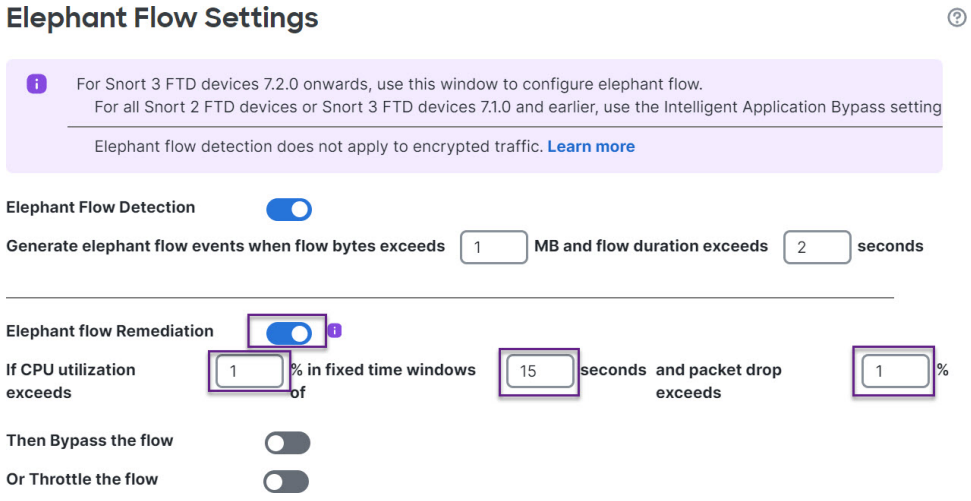

Step 7 Enable the **Bypass the flow** toggle button and click the **Select Applications/Filters** radio button.

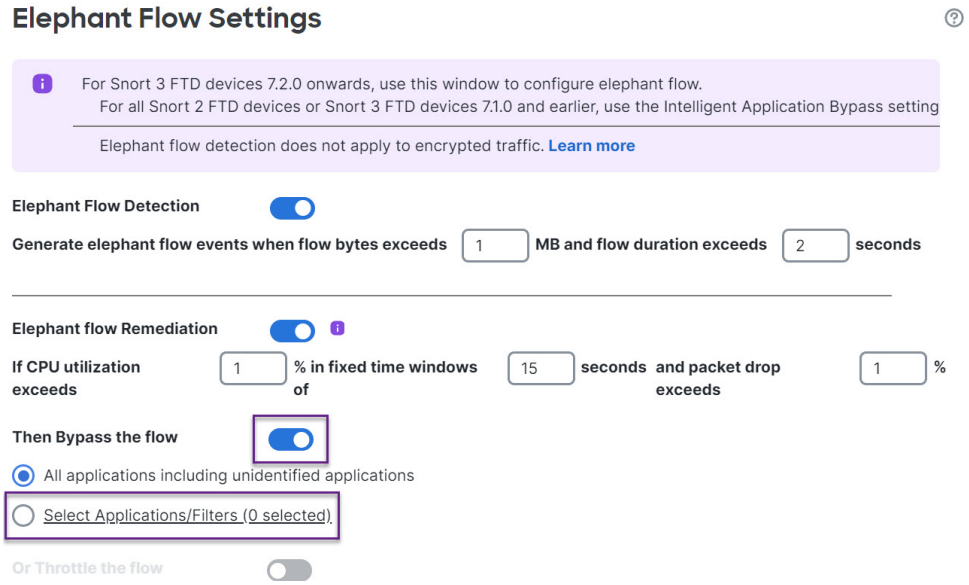

Step 8 Under **Application Filters**, search for and select the **WebEx** application, add it to the rule, and click **Save**. This means that WebEx connections are trusted and prioritized and will skip Snort inspection if these WebEx connections are detected as elephant flows, based on the configured parameters.

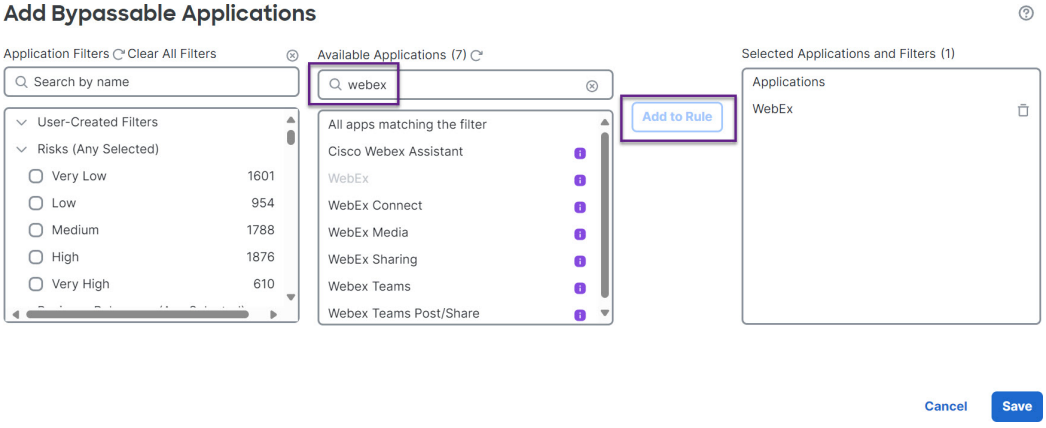

Step 9 Enable the **Throttle** toggle button to throttle the remaining flows (causing duress). This ensures that all the other flows are slowed down in 10 percent increments until the Snort duress condition is met.

Step 10 Click **OK**.

Step 11 Click **Save**.

What to do next

Deploy configuration changes. See Deploy Configuration Changes, on page 2055.

View Events for Elephant Flows

After configuring your elephant flow settings, monitor your connection events to see if any flows are detected, bypassed, or throttled. You can see this information in the **Reason** field of your connection events. The three types for elephant flow connections are:

- Elephant Flow
- Elephant Flow Throttled
- Elephant Flow Trusted

Procedure

Step 1 Choose **Analysis** > **Connections**. You can also view the events from the **Unified Events** viewer.

Step 2 In the **Connection Events** page, from the **Predefined Search** drop-down list, choose **Elephant Flows** to display elephant flow events.

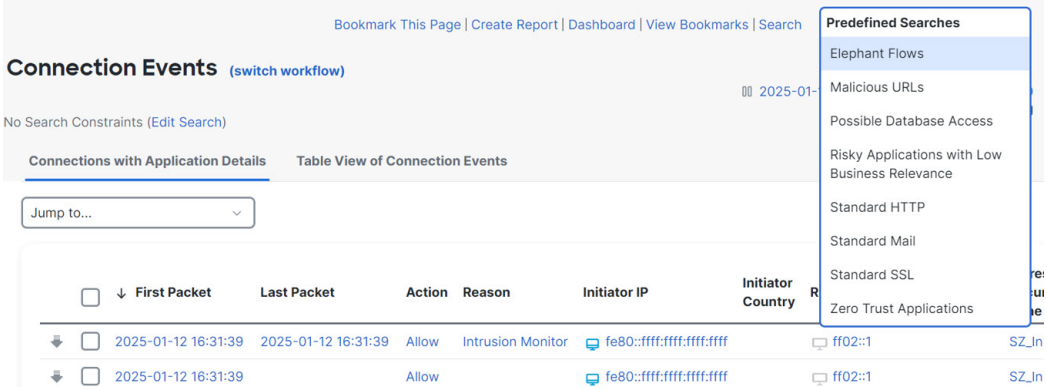

Tip
To see **Elephant Flow Trusted** or **Elephant Flow Throttled** event types, click the **Edit Search** link on the top-left corner of the page and in the **Reason** field, choose **Elephant Flows** in the left panel. Enter **Elephant Flow Trusted** or **Elephant Flow Throttled**, depending on what you want to search.

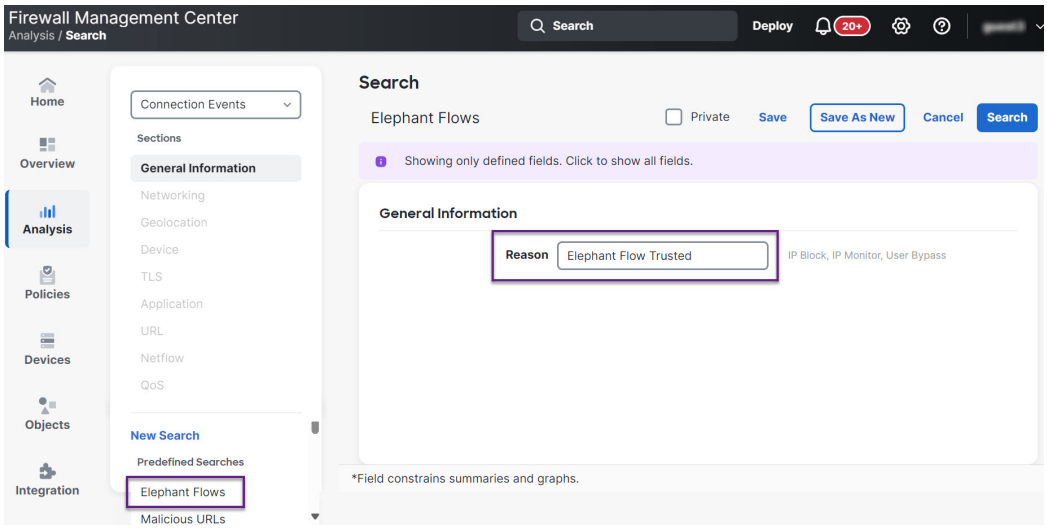

Step 3 View the elephant flow that was detected mid-flow and the **Reason** field shows **Elephant Flow**. At the end of the flow, it was bypassed and the **Reason** field shows **Elephant Flow Trusted**.

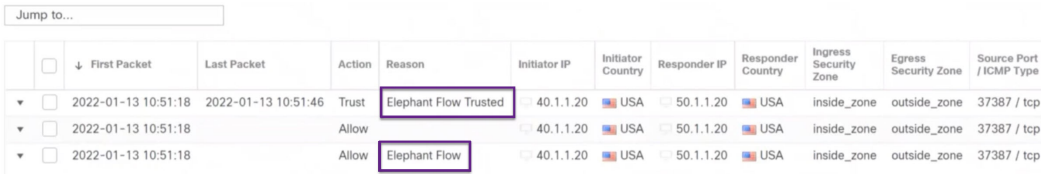

Configure Elephant Flow Remediation Exemption

You can configure L4 access control list (ACL) rules for flows that must be exempted from remediation. If a flow is detected as an elephant flow and it matches the rules that are defined, that flow is exempted from the remediation action.

Before you begin

You must be running management center 7.4.0 or later and the managed threat defense must also be 7.4.0 or later.

Procedure

Step 1 Choose **Policies** > **Access Control**.

Step 2 Click **Edit** (✎) next to the access control policy you want to edit.

Step 3 Choose **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.

Step 4 Click **Edit** (✎) next to **Elephant Flow Settings**.

Step 5 Ensure that you have configured the elephant flow detection and remediation parameters. See Configure Elephant Flow Parameters, on page 2189.

Step 6 Click the **Add Rule** button next to **Remediation Exemption Rules**.

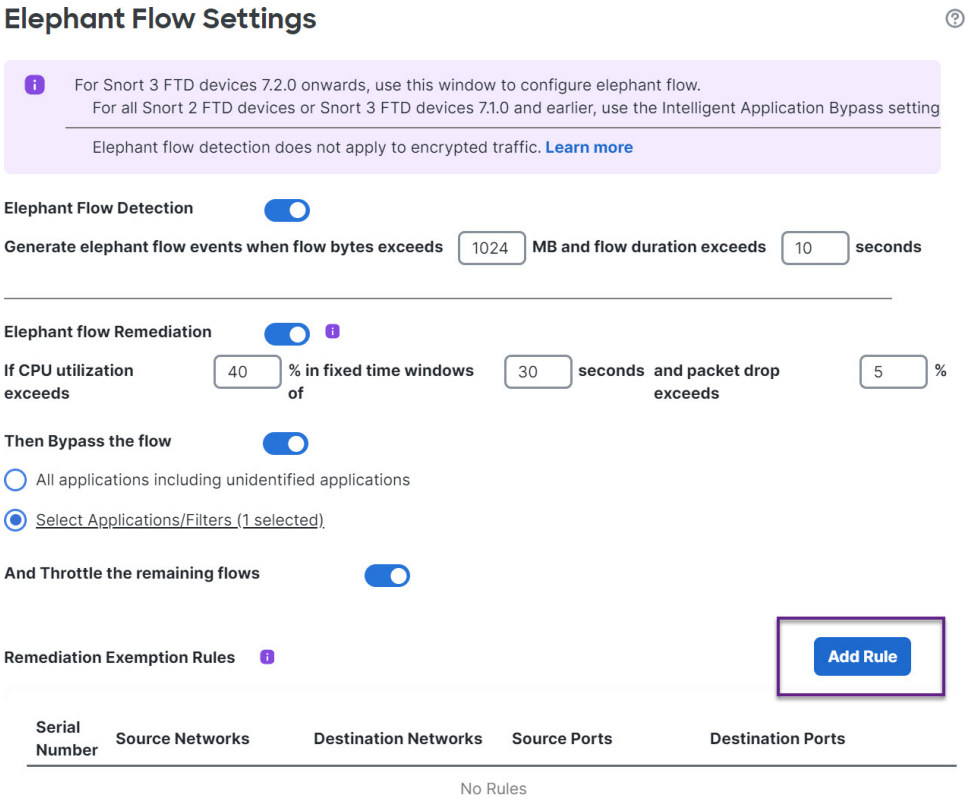

Step 7 From the list of **Available Networks**, choose the configured host to exempt from elephant flow remediation. For the purposes of this example, we have created a host called "Host1_Exception."

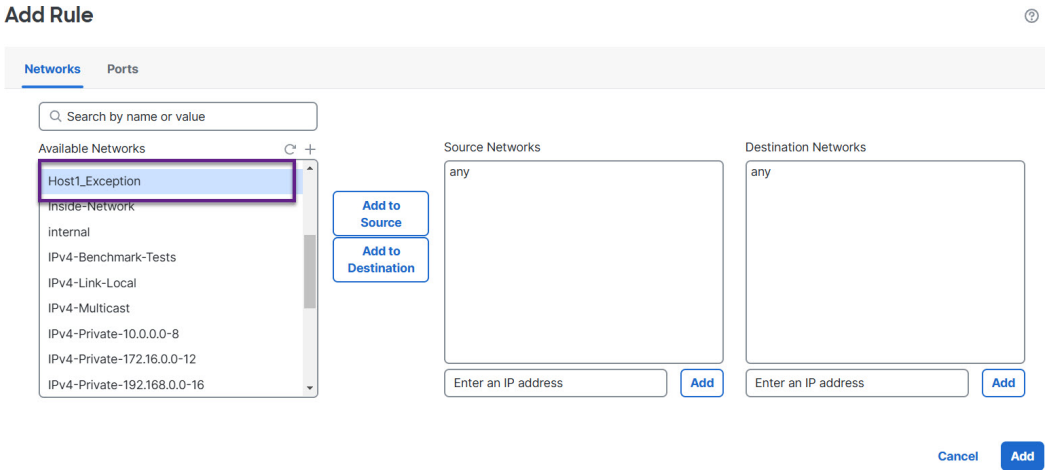

Step 8 Click **Add to Source** or **Add to Destination** (as required) to add this host to the source or destination.

Step 9 Click the **Ports** tab.

Step 10 For the source port, choose **Protocol** as TCP and enter **80** as the destination port, and click **Add**.

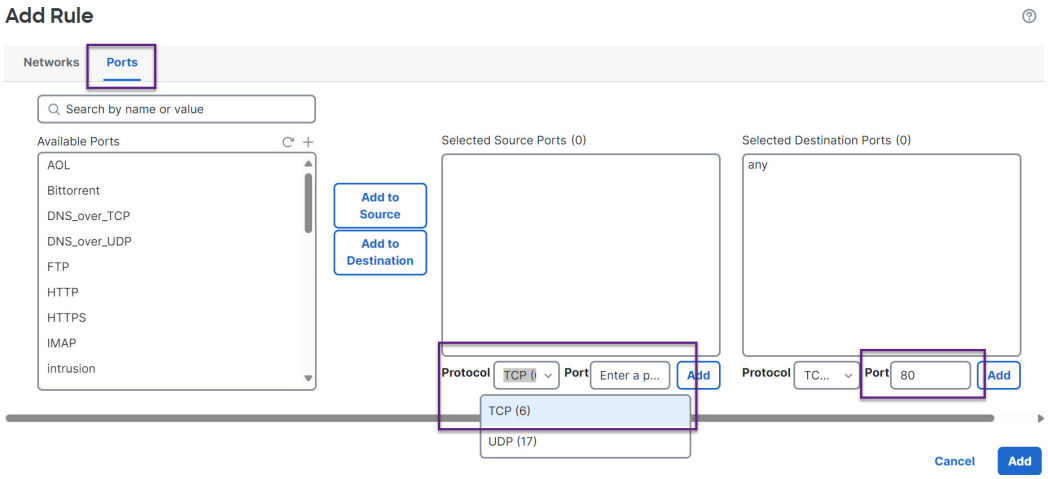

Step 11 Click **OK**.

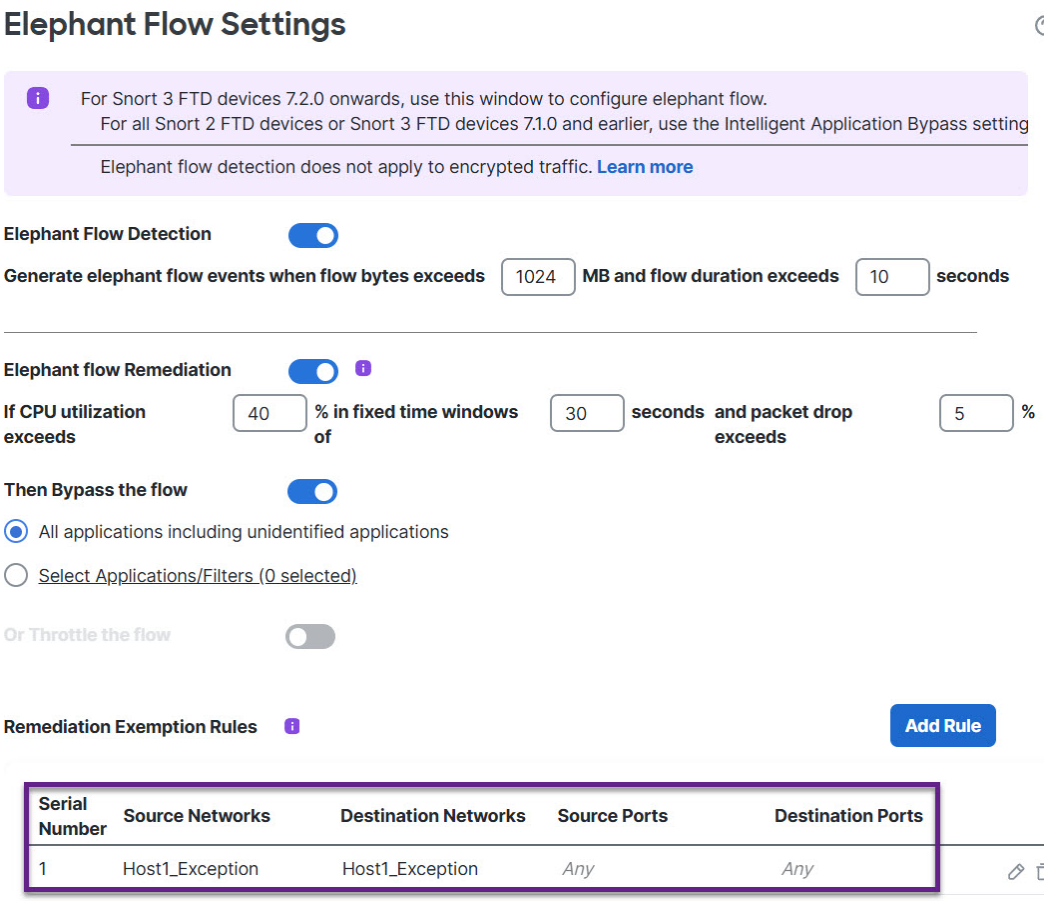

Step 12 Click **Save**.

What to do next

Deploy configuration changes. See Deploy Configuration Changes, on page 2055.

View Events for Elephant Flow Remediation Exemption

Procedure

Step 1 Choose **Analysis** > **Connections**. You can also view the events from the **Unified Events** viewer.

Step 2 View the elephant flows that were exempted from remediation. The **Reason** field shows **Elephant Flow Exempted**.

Additional References

For detailed conceptual information, see the Elephant Flow Detection for Snort 3 chapter in this guide or the content in the following link:

- Elephant Flow Detection

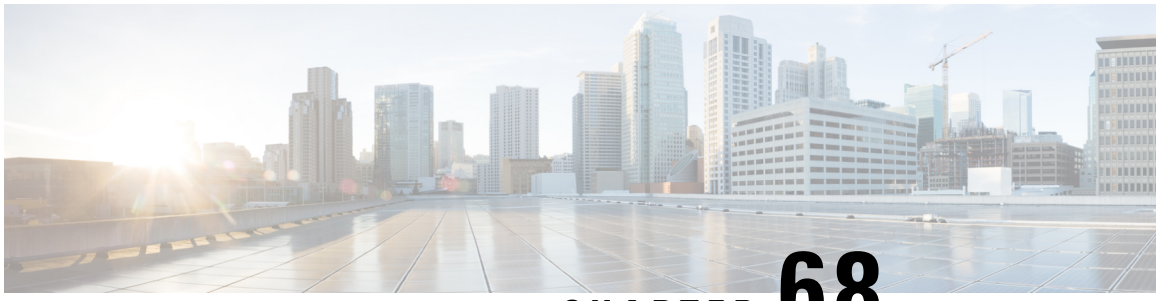

CHAPTER 68

Mitigate Threats Using MITRE Framework in Snort 3 Intrusion Policies

- About MITRE ATT&CK Framework, on page 2197
- Benefits of MITRE Framework, on page 2198
- Sample Business Scenario for MITRE Network, on page 2198
- Prerequisites for MITRE Framework, on page 2198
- View and Edit Your Snort 3 Intrusion Policy, on page 2198
- View Intrusion Events, on page 2204
- Additional References, on page 2206

About MITRE ATT&CK Framework

The MITRE ATT&CK Framework is a comprehensive knowledge base that outlines the tactics, techniques, and procedures (TTPs) used by threat actors to compromise systems. It organizes these TTPs into matrices for different operating systems and platforms, mapping each attack stage (tactics) to specific methods (techniques). Each technique includes information about execution, procedures, defenses, detections, and real-world examples.

Note See https://attack.mitre.org for additional information about MITRE ATT&CK.

The management center uses the MITRE ATT&CK Framework to enhance threat detection and response, incorporating the following capabilities:

- Intrusion events include TTPs, allowing administrators to manage traffic with greater granularity by grouping rules according to vulnerability type, target system, or threat category.
- Select malware events use TTPs, enhancing the ability to detect and respond to threats.
- Unified and Classic Event Viewers display tactics, techniques, attack lifecycle graphs, and contextual enrichments from the Talos taxonomy. These enrichments include MITRE tags and a hierarchical view of associated tactics, techniques, and sub-techniques. You can also filter events using MITRE identifiers.

Benefits of MITRE Framework

- MITRE Tactics, Techniques, and Procedures (TTPs) are added to intrusion events, which enable administrators to act on traffic, based on the MITRE ATT&CK framework. This enables administrators to view and handle traffic with more granularity, and group rules by vulnerability type, target system, or threat category.
- You can organize intrusion rules according to the MITRE ATT&CK framework. This allows you to customize policies according to specific attacker tactics and techniques.

Sample Business Scenario for MITRE Network

A large corporate network uses Snort 3 as its primary intrusion detection and prevention system. In a rapidly evolving threat landscape, adoption of robust network security measures is necessary and important. Network administrators need to know if the configured policies are finding traffic of interest and if they are tracking a known attack group. For example, you may want to know if adversaries are attempting to take advantage of a weakness in your systems or applications in order to cause unanticipated behavior. The weakness in the system can be a bug, a glitch, or a design vulnerability. The applications may be websites, databases, standard services, such as Server Message Block (SMB) or Secure Shell (SSH), network device administration and management protocols or applications, such as web servers and related services.

The insights provided by the MITRE framework provides administrators with a more precise opportunity to specify protection for specific assets and protect their network from specific threat groups.

Prerequisites for MITRE Framework

- You must be running Secure Firewall Management Center and Secure Firewall Threat Defence Version 7.3.0 or later with Snort 3.
- You must have at least one intrusion policy. See Create a Custom Snort 3 Intrusion Policy , on page 2060.

View and Edit Your Snort 3 Intrusion Policy

Procedure

Step 1 Choose **Policies** > **Intrusion**.

Step 2 Ensure that the **Intrusion Policies** tab is chosen.

Step 3 Click **Snort 3 Version** next to the intrusion policy that you want to view or edit.

Step 4 Close the Snort helper guide that is displayed.

Step 5 Click the **Group Overrides** layer.

This layer lists all the categories of rule groups in an hierarchical structure. You can drill down to the last leaf rule group under each rule group.

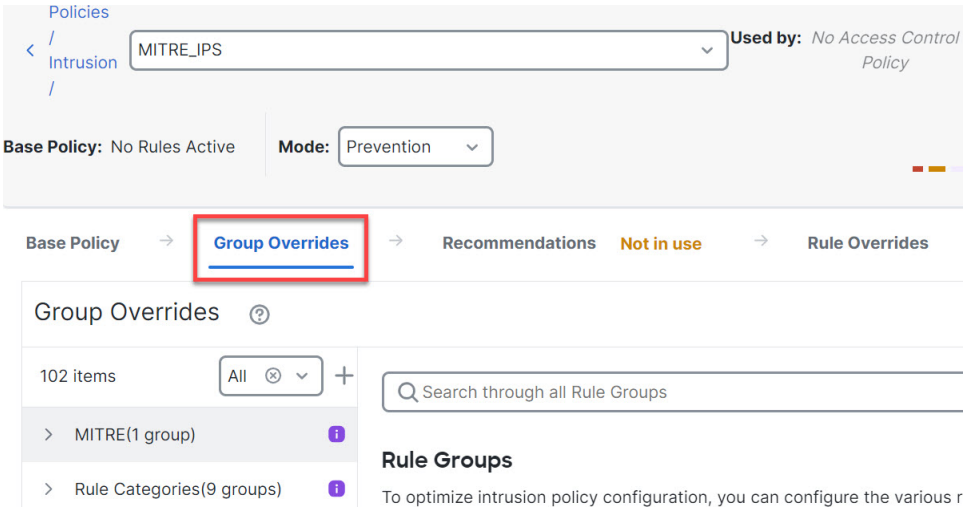

Step 6 Under **Group Overrides**, ensure that **All** is chosen in the drop-down list, so that all the rule groups for the corresponding intrusion policy are visible in the left pane.

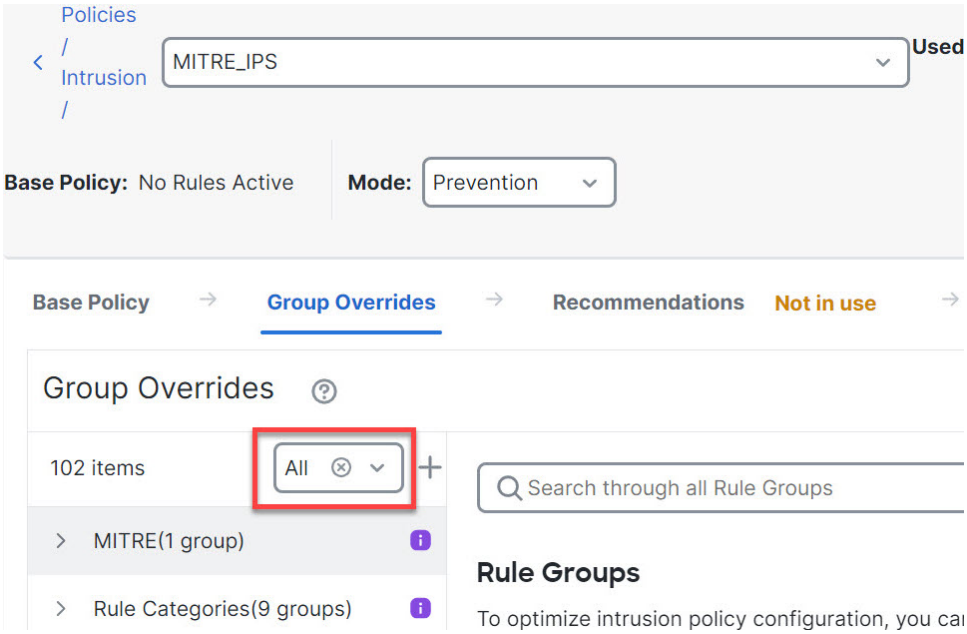

Step 7 Click **MITRE** in the left pane.

Note
Depending on your specific requirements, you can choose the **Rule Categories** rule group or any other rule group and subrule groups under it. All the rule groups use the MITRE framework.

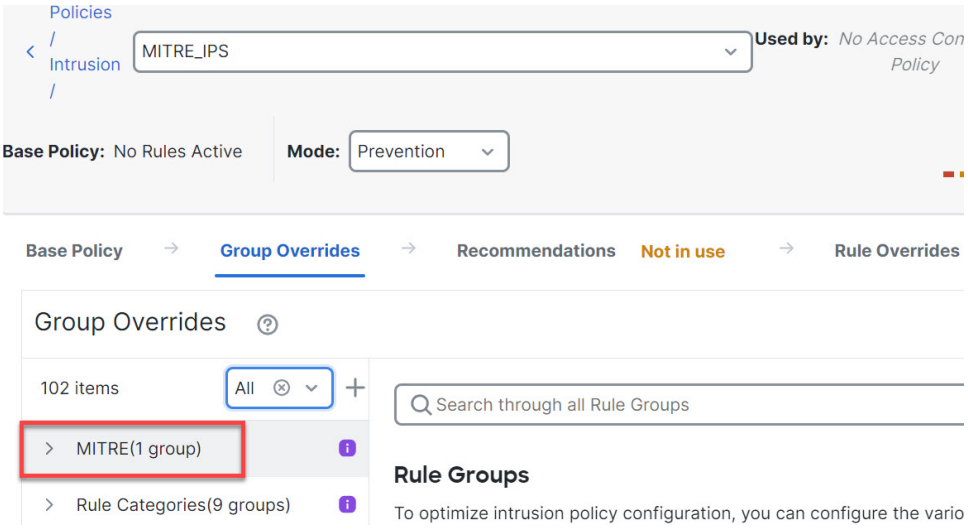

Step 8 Under **MITRE**, click **ATT&CK Framework** to drill down.

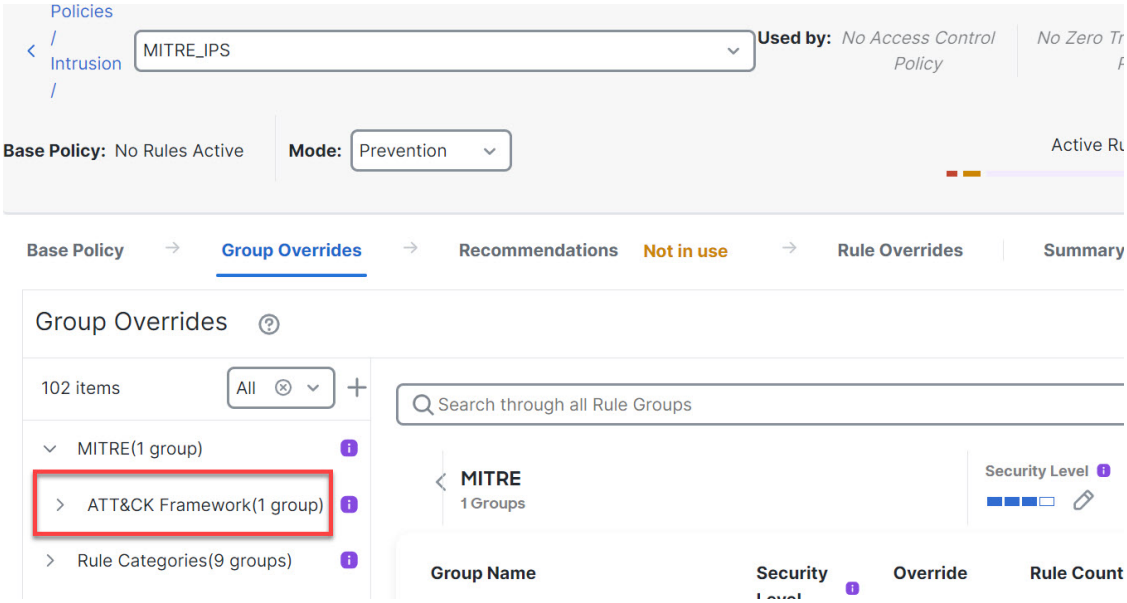

Step 9 Under **ATT&CK Framework**, click **Enterprise** to expand it.

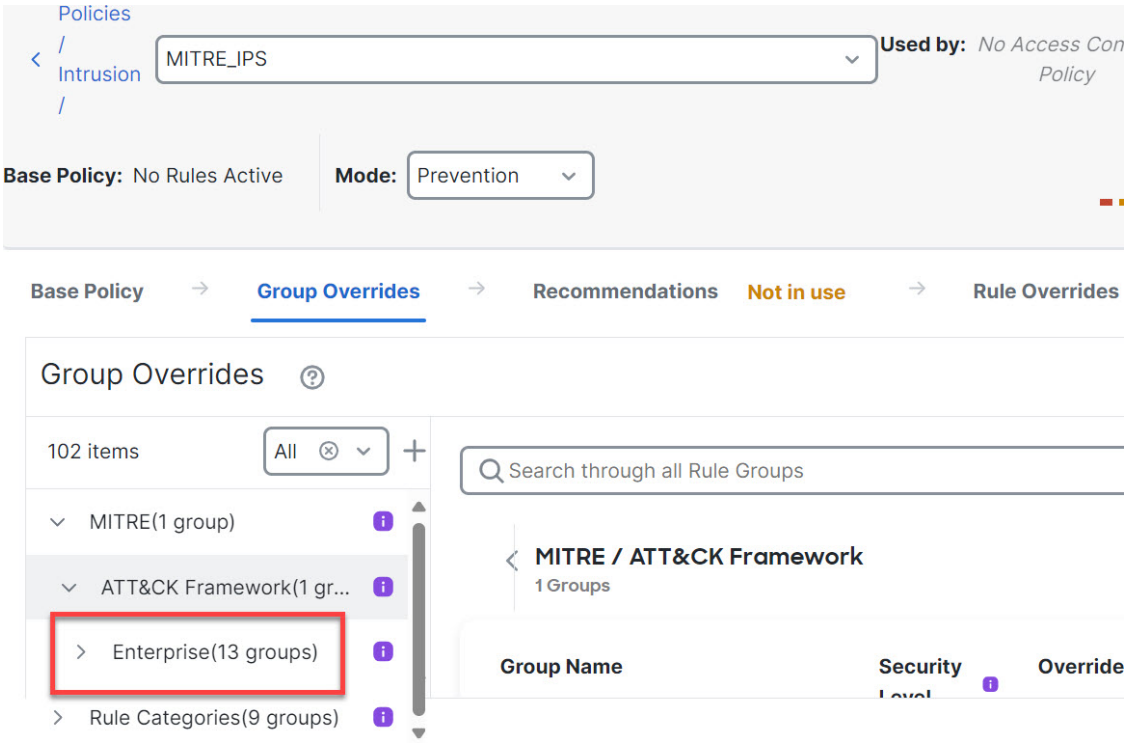

Step 10 Click the **Edit** (✎) icon next to the **Security Level** of the rule group to make bulk changes to the security level of all the associated rule groups under the **Enterprise** rule group category.

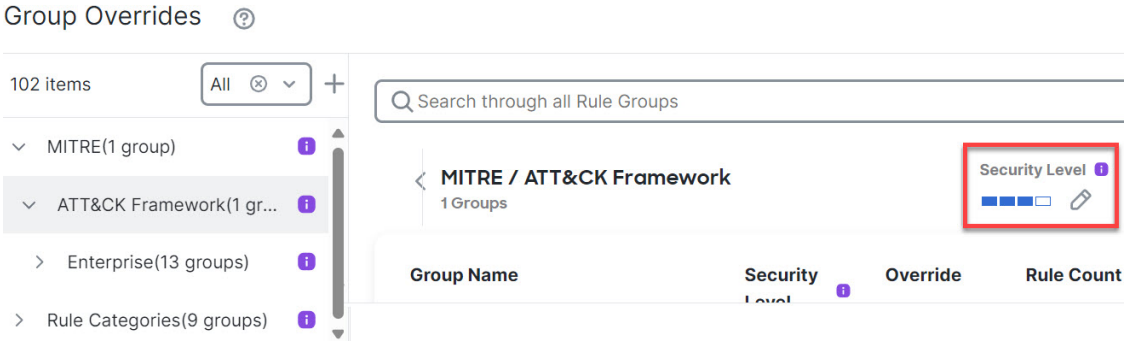

Step 11 In the **Edit Security Level** window, choose a **Security Level** (in this example, **3**), and click **Save**.

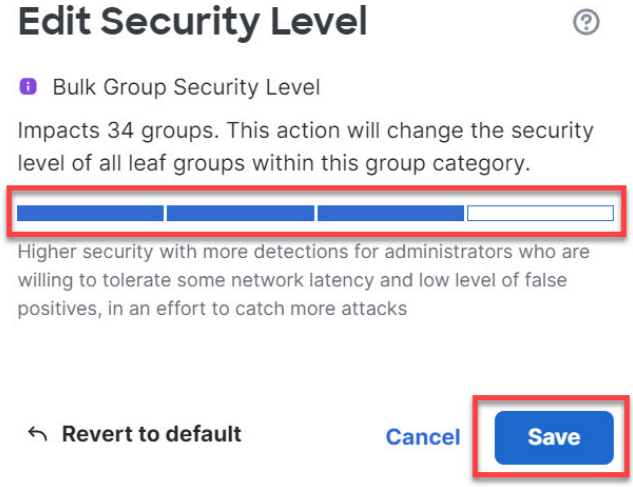

Step 12 Under **Enterprise**, click **Initial Access** to expand it.

Step 13 Under **Initial Access**, click **Exploit Public-Facing Application**, which is the last leaf group.

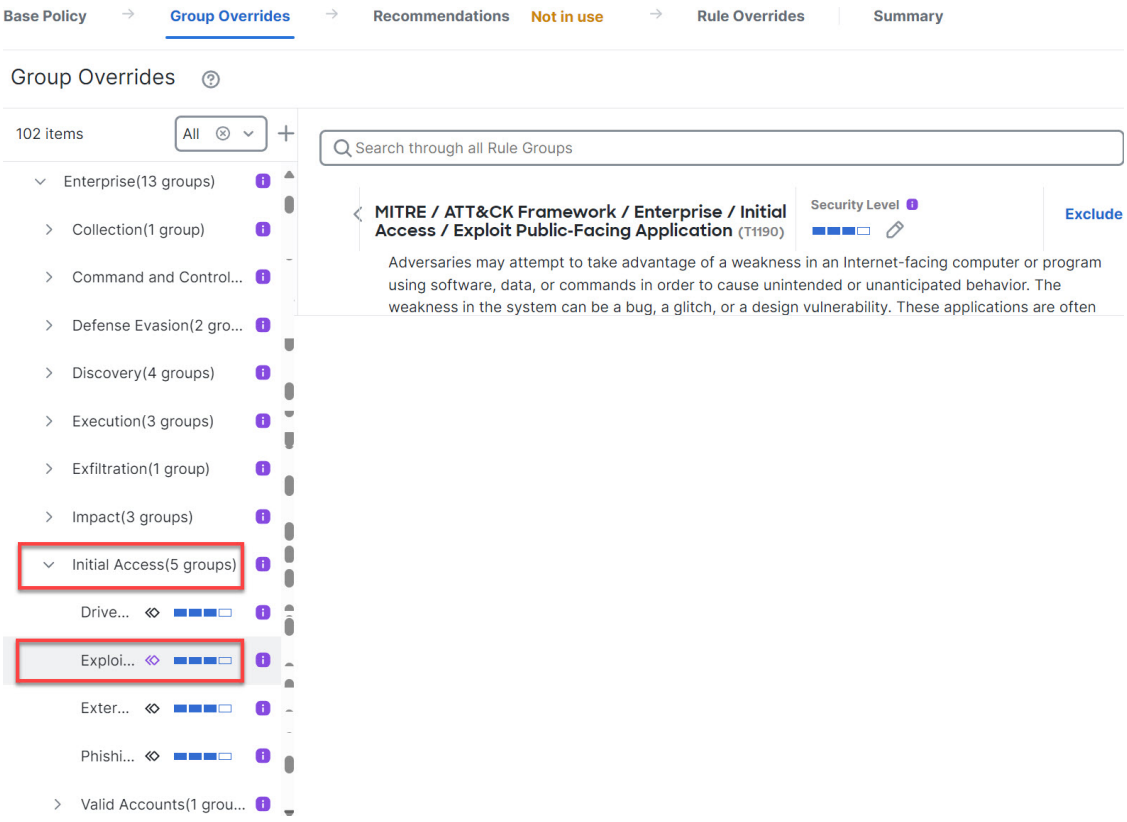

Step 14 Click **View Rules in Rule Overrides** to view the different rules, rule details, rule actions, and so on, for the different rules. You can change the rule actions for one or multiple rules in the **Rule Overrides** layer.

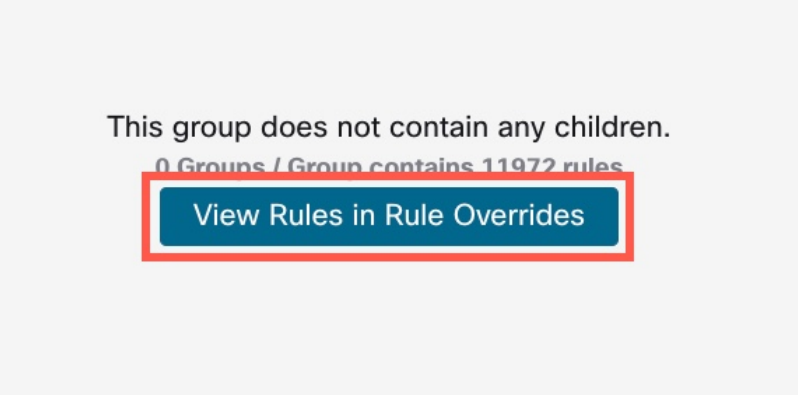

Step 15 Click the **Recommendations** layer and then click **Start** to start using Cisco-recommended rules. You can use the intrusion rule recommendations to target the vulnerabilities that are associated with the host assets detected in the network. For more information, see Generate New Secure Firewall Recommendations in Snort 3, on page 2095.

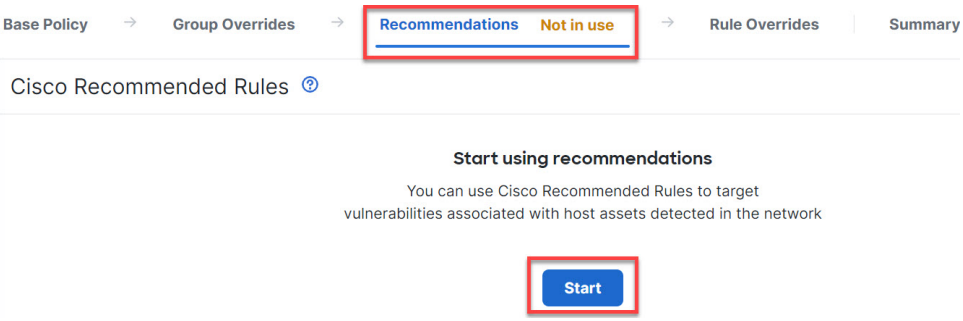

Step 16 Click the **Summary** layer for a holistic view of the current changes to the policy. Based on the rule overrides, security-level changes, and generation of Cisco-recommended rules, you can view the rule distribution of the policy, group overrides, rule overrides, rule recommendations, and so on, to verify your changes.

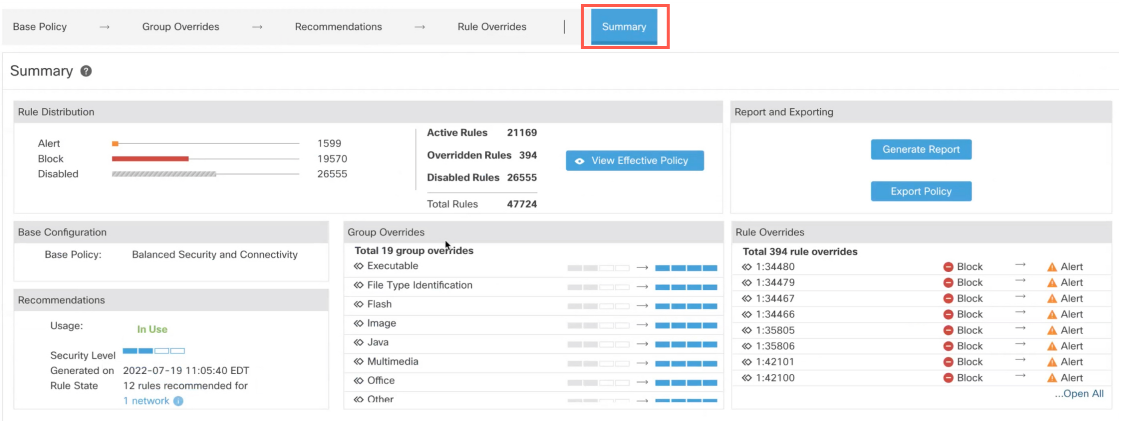

What to do next

Deploy your intrusion policy to detect and log events that are triggered by the Snort rules. See Deploy Configuration Changes, on page 2055.

View Intrusion Events

You can view the MITRE ATT&CK techniques and rule groups in the intrusion events on the **Classic Event Viewer** and **Unified Event Viewer** pages. Talos provides mappings from Snort rules (GID:SID) to MITRE ATT&CK techniques and rule groups. These mappings are installed as part of the Lightweight Security Package (LSP).

Procedure

Step 1 Click **Analysis** and select **Events** under **Intrusions**.

Step 2 Click the **Table View of Events** tab.

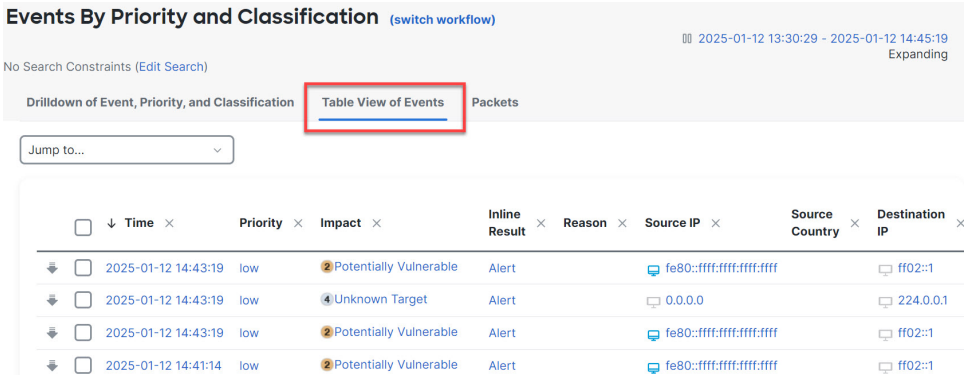

Step 3 Under **MITRE ATT&CK**, you can see the techniques for an intrusion event. Click **1 Technique** to view the MITRE ATT&CK techniques.

In this example, **Exploit Public-Facing Application** is the technique.

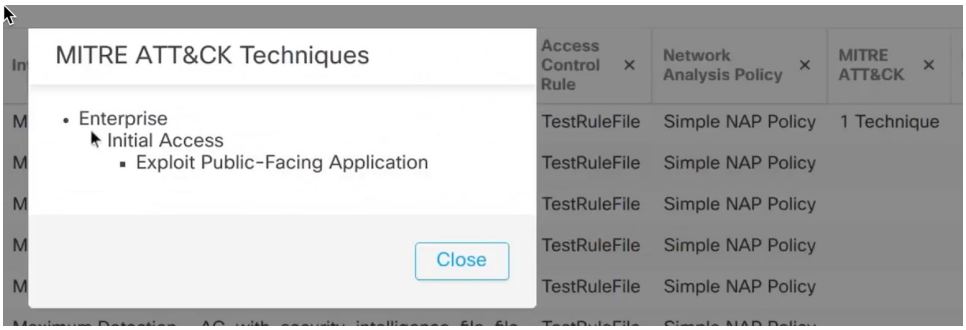

Step 4 Click **Close**.

Step 5 Click **Analysis** and select **Unified Events**.

Step 6 If not enabled, click the column selector icon to enable the **MITRE ATT&CK** and **Rule Group** columns.

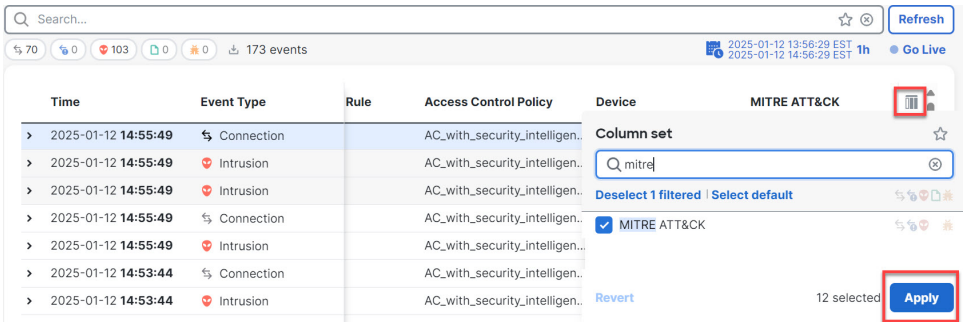

Step 7 In this example, the intrusion event is triggered by an event that is mapped to one rule group. Click **1 Group** under the **Rule Group** column.

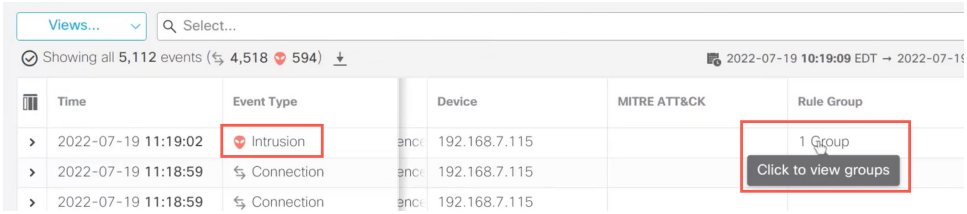

Step 8 You can view **Protocol**, which is the parent rule group, and the DNS rule group under it. Choose **Protocol** > **DNS** to search for all the intrusion events that have at least one rule group that is .

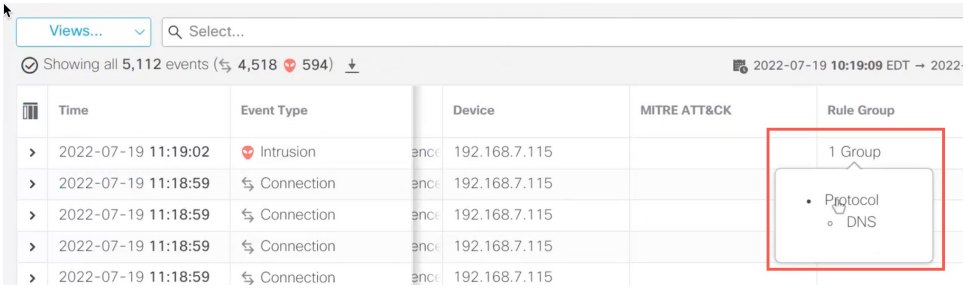

The search results are displayed.

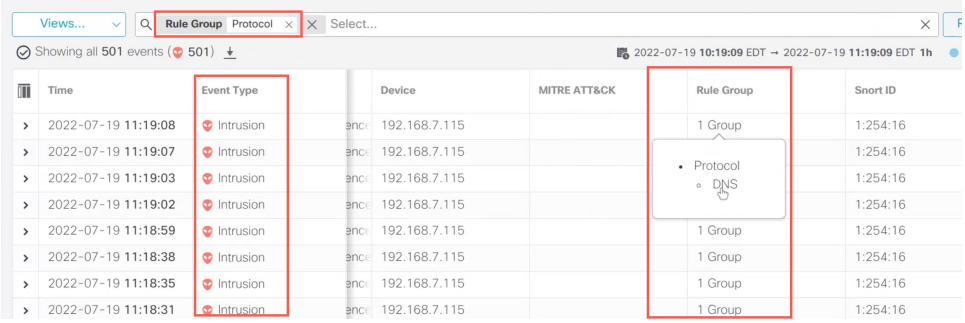

Additional References

- Intrusion Policy in Snort 3
- Edit Snort 3 Intrusion Policies, on page 2061
- MITRE Information in Malware Events

PART X

Network Malware Protection and File Policies

- Network Malware Protection and File Policies, on page 2209

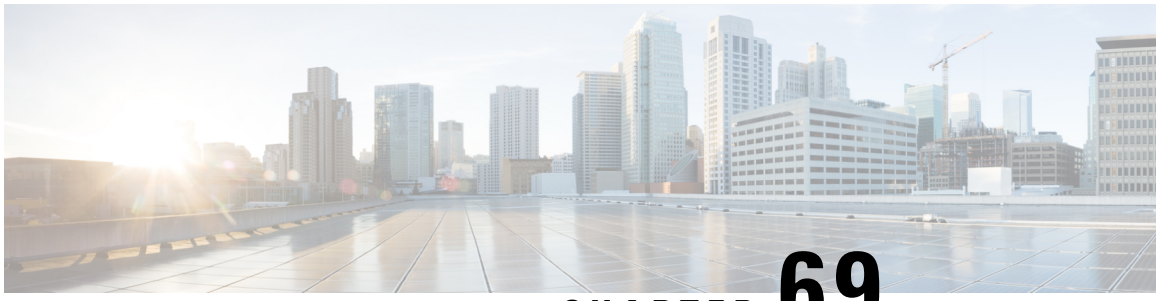

CHAPTER 69

Network Malware Protection and File Policies

The following topics provide an overview of file control, file policies, file rules, Advanced Malware Protection (AMP), cloud connections, and dynamic analysis connections.

- About Network Malware Protection and File Policies, on page 2209
- Requirements and Prerequisites for File Policies, on page 2210
- License Requirements for File and Malware Policies, on page 2211
- Best Practices for File Policies and Malware Detection , on page 2211
- How to Configure Malware Protection, on page 2214
- Cloud Connections for Malware Protection, on page 2219
- File Policies and File Rules, on page 2228
- Retrospective Disposition Changes, on page 2243
- File and Malware Inspection Performance and Storage Options, on page 2243
- Tuning File and Malware Inspection Performance and Storage, on page 2244
- (Optional) Malware Protection with Secure Endpoint, on page 2245
- History for Network Malware Protection and File Policies, on page 2250

About Network Malware Protection and File Policies

To detect and block malware, use file policies. You can also use file policies to detect and control traffic by file type.

Advanced Malware Protection (AMP) for Firepower can detect, capture, track, analyze, log, and optionally block the transmission of malware in network traffic. In the Secure Firewall Management Center web interface, this feature is called *malware defense*, formerly called *AMP for Firepower*. Advanced Malware Protection identifies malware using managed devices deployed inline and threat data from the Cisco cloud.

You associate file policies with access control rules that handle network traffic as part of your overall access control configuration.

When the system detects malware on your network, it generates file and malware events. To analyze file and malware event data, see the *File/Malware Events and Network File Trajectory* chapter in the Cisco Secure Firewall Management Center Administration Guide .

File Policies

A file policy is a set of configurations that the system uses to perform malware protection and file control, as part of your overall access control configuration. This association ensures that before the system passes a file in traffic that matches an access control rule's conditions, it first inspects the file. Consider the following diagram of a simple access control policy in an inline deployment.

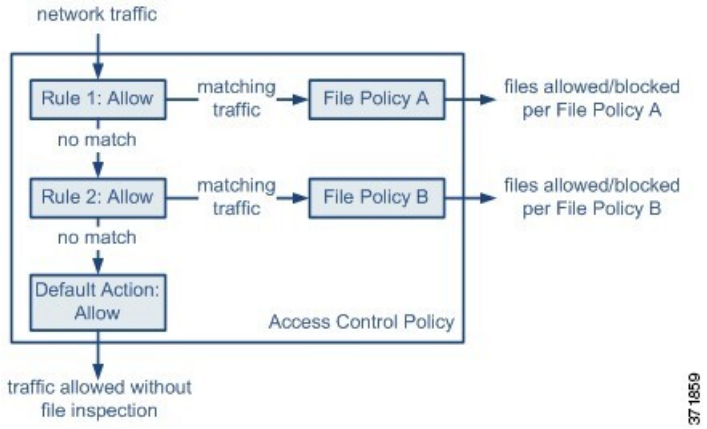

The policy has two access control rules, both of which use the Allow action and are associated with file policies. The policy's default action is also to allow traffic, but without file policy inspection. In this scenario, traffic is handled as follows:

- Traffic that matches `Rule 1` is inspected by `File Policy A`.
- Traffic that does not match `Rule 1` is evaluated against `Rule 2`. Traffic that matches `Rule 2` is inspected by `File Policy B`.
- Traffic that does not match either rule is allowed; you cannot associate a file policy with the default action.

By associating different file policies with different access control rules, you have granular control over how you identify and block files transmitted on your network.

Requirements and Prerequisites for File Policies

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Access Admin

License Requirements for File and Malware Policies

To Do This	License Required	File Rule Action
Block or allow all files of a particular type (for example, block all .exe files)	IPS (for Firewall Threat Defense devices) Protection (for Classic devices)	Allow, Block, Block with Reset
Selectively allow or block files based on a judgment that it contains or is likely to contain malware	IPS (for Firewall Threat Defense devices) Protection (for Classic devices) Malware Defense	Malware Cloud Lookup, Block Malware
Store files	IPS (for Firewall Threat Defense devices) Protection (for Classic devices) Malware Defense	Any file rule action with **Store Files** selected

For details about Malware Defense licenses, see:

- *Malware Defense Licenses* in the Cisco Secure Firewall Management Center Administration Guide

Best Practices for File Policies and Malware Detection

In addition to the items described below, follow the steps in How to Configure Malware Protection, on page 2214 and referenced topics.

File Rule Best Practices

Note the following guidelines and limitations when configuring file rules:

- A rule configured to block files in a passive deployment does not block matching files. Because the connection continues to transmit the file, if you configure the rule to log the beginning of the connection, you may see multiple events logged for this connection.
- A policy can include multiple rules. When you create the rules, ensure that no rule is "shadowed" by a previous rule.
- The file types supported for dynamic analysis are a subset of the file types supported for other types of analysis. To view the file types supported for each type of analysis, navigate to the file rule configuration page, select the **Block Malware** action, and select the checkboxes of interest.

To ensure that the system examines all file types, create separate rules (within the same policy) for dynamic analysis and for other types of analysis.

- If a file rule is configured with a **Malware Cloud Lookup** or **Block Malware** action and the Firewall Management Center cannot establish connectivity with the AMP cloud, the system cannot perform any configured rule action options until connectivity is restored.
- Cisco recommends that you enable **Reset Connection** for the **Block Files** and **Block Malware** actions to prevent blocked application sessions from remaining open until the TCP connection resets. If you do not reset connections, the client session will remain open until the TCP connection resets itself.
- If you are monitoring high volumes of traffic, do **not** store all captured files, or submit all captured files for dynamic analysis. Doing so can negatively impact system performance.
- You cannot perform malware analysis on all file types detected by the system. After you select values from the **Application Protocol**, **Direction of Transfer**, and **Action** drop-down lists, the system constrains the list of file types.

File Detection Best Practices

Consider the following notes and limitations for file detection:

- If adaptive profiling is not enabled, access control rules cannot perform file control, including AMP.
- If a file matches a rule with an application protocol condition, file event generation occurs after the system successfully identifies a file's application protocol. Unidentified files do not generate file events.
- FTP transfers commands and data over different channels. In a passive or inline tap mode deployment, the traffic from an FTP data session and its control session may not be load-balanced to the same internal resource.
- If the total number of bytes for all file names for files in a POP3, POP, SMTP, or IMAP session exceeds 1024, file events from the session may not reflect the correct file names for files that were detected after the file name buffer filled.
- When transmitting text-based files over SMTP, some mail clients convert newlines to the CRLF newline character standard. Since Mac-based hosts use the carriage return (CR) character and Unix/Linux-based hosts use the line feed (LF) character, newline conversion by the mail client can modify the size of the file. Note that some mail clients default to newline conversion when processing an unrecognizable file type.
- To detect ISO files, set the "Limit the number of bytes inspected when doing file type detection" option to a value greater than 36870, as described in File and Malware Inspection Performance and Storage Options, on page 2243.
- .Exe files inside some .rar archives cannot be detected, including possibly rar5.
- If the file disposition is Neutral, it is an unknown disposition.

File Blocking Best Practices

Consider the following notes and limitations for file blocking:

- If an end-of-file marker is not detected for a file, regardless of transfer protocol, the file will not be blocked by a **Block Malware** rule or the custom detection list. The system waits to block the file until the entire file has been received, as indicated by the end-of-file marker, and blocks the file after the marker is detected.
- If the end-of-file marker for an FTP file transfer is transmitted separately from the final data segment, the marker will be blocked and the FTP client will indicate that the file transfer failed, but the file will actually completely transfer to disk.
- File rules with **Block Files** and **Block Malware** actions block automatic resumption of file download via HTTP by blocking new sessions with the same file, URL, server, and client application detected for 24 hours after the initial file transfer attempt occurs.
- In rare cases, if traffic from an HTTP upload session is out of order, the system cannot reassemble the traffic correctly and therefore will not block it or generate a file event.
- If you transfer a file over NetBIOS-ssn (such as an SMB file transfer) that is blocked with a **Block Files** rule, you may see a file on the destination host. However, the file is unusable because it is blocked after the download starts, resulting in an incomplete file transfer.
- If you create file rules to detect or block files transferred over NetBIOS-ssn (such as an SMB file transfer), the system does not inspect the ongoing file transfers. However, the system inspects the new files that are transferred after you deploy an access control policy invoking the file policy.
- SMB has a functionality called multi-channel which creates multiple parallel sessions with the same IP address and different ports. For a given transaction over multi-channel, the file download is multiplexed across these sessions which is not inspected by the system as a single file.
- Files transferred concurrently in a single TCP or SMB session are not inspected.
- In a cluster environment, if an existing SMB session is moved to a new device due to a cluster role change or a device failure, then the files in any ongoing file transfers may not be inspected.
- Some SMB file transfers between Microsoft Windows systems use very high TCP window size for quick file transfers. To detect or block such file transfers, it is recommended that you increase the value of **Maximum Queued Bytes** and **Maximum Queued Segments** under **Network Analysis Policy** > **TCP Stream Configuration** > **Troubleshooting Options**.
- If you configure Firewall Threat Defense high availability, and failover occurs while the original active device is identifying the file, the file type is not synced. Even if your file policy blocks that file type, the new active device downloads the file.

File Policy Best Practices

Note the following general guidelines and limitations when configuring file policies.

- You can associate a single file policy with an access control rule whose action is **Allow**, **Interactive Block**, or **Interactive Block with reset**.
- You **cannot** use a file policy to inspect traffic handled by the access control default action.
- For a new policy, the web interface indicates that the policy is not in use. If you are editing an in-use file policy, the web interface tells you how many access control policies use the file policy. In either case, you can click the text to jump to the Access Control Policies page.

- For file blocking to work, the NAP policy you apply to the access control policy must be operating in Protection mode, also known as Inline mode.
- Based on your configuration, you can either inspect a file the first time the system detects it, and wait for a cloud lookup result, or pass the file on this first detection without waiting for the cloud lookup result.
- By default, file inspection of encrypted payloads is disabled. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has file inspection configured.

Attention The File Inpsect preprocessor with the following generator IDs (GIDs) are enabled by default for file/malware policy: GID: 146 and GID: 147.

- If you enable an access control policy with a file policy that uses either a Malware action or a Store Files option, the computing power and system performance of the device is reduced.

How to Configure Malware Protection

This topic summarizes the steps you must take to set up your system to protect your network from malicious software.

Procedure

Step 1 Plan and Prepare for Malware Protection, on page 2215

Step 2 Configure File Policies, on page 2216

Step 3 Add File Policies to Your Access Control Configuration, on page 2216

Step 4 Configure network discovery policies to associate file and malware events with hosts on your network.

(Do not simply turn on network discovery; you must configure it to discover hosts on your network to build a network map of your organization.)

See Network Discovery Policies, on page 2643 and subtopics.

Step 5 Deploy policies to managed devices.

See Deploy Configuration Changes, on page 259.

Step 6 Test your system to be sure it is processing malicious files as you expect it to.

Step 7 Set Up Maintenance and Monitoring of Malware Protection, on page 2218

What to do next

- (Optional) To further enhance detection of malware in your network, deploy and integrate Cisco's Secure Endpoint product. See (Optional) Malware Protection with Secure Endpoint, on page 2245 and subtopics.

- Understand how to investigate file and malware events.

  See *File/Malware Events and Network File Trajectory* in the Cisco Secure Firewall Management Center Administration Guide.

Plan and Prepare for Malware Protection

This procedure is the first set of steps in the complete process for configuring your system to provide malware protection.

Procedure

Step 1 Purchase and install licenses.

See License Requirements for File and Malware Policies, on page 2211 and *Licenses* in the Cisco Secure Firewall Management Center Administration Guide.

Step 2 Understand how file policies and malware protection fit into your access control plan.

See the chapter Access Control Overview, on page 1723.

Step 3 Understand the file analysis and malware protection tools.

See File Rule Actions, on page 2234 and subtopics.

Consider also Advanced and Archive File Inspection Options, on page 2229.

Step 4 Determine whether you will use public clouds or private (on-premises) clouds for malware protection (file analysis and dynamic analysis.)

See Cloud Connections for Malware Protection, on page 2219 and subtopics.

Step 5 If you will use private (on-premises) clouds for malware protection: Purchase, deploy, and test those products.

For information, contact your Cisco sales representative or authorized reseller.

Step 6 Configure your firewall to allow communications with your chosen clouds.

See *Security, Internet Access, and Communication Ports* in the Cisco Secure Firewall Management Center Administration Guide.

Step 7 Configure connections between Firepower and the malware protection clouds (public or private, as applicable).

- For the AMP cloud, see Change AMP Options, on page 2224.
- If you have deployed an on-premises Secure Malware Analytics Appliance, see Connect to an On-Premises Dynamic Analysis Appliance, on page 2225. (Access to the public Secure Malware Analytics Cloud does not require configuration.)

What to do next

Continue with the next step in the malware protection workflow:

See How to Configure Malware Protection, on page 2214.

Configure File Policies

Before you begin

Complete the tasks up to this point in the malware protection workflow:

See How to Configure Malware Protection, on page 2214.

Procedure

Step 1 Review file policy and file rule restrictions.

See Best Practices for File Policies and Malware Detection , on page 2211 and subtopics.

Step 2 Create a file policy.

See Create or Edit a File Policy, on page 2228.

Step 3 Create rules within your file policy.

See File Rules, on page 2233 and subtopics.

Step 4 Configure advanced options.

See Advanced and Archive File Inspection Options, on page 2229.

What to do next

Continue with the next step in the malware protection workflow:

See How to Configure Malware Protection, on page 2214.

Add File Policies to Your Access Control Configuration

An access control policy can have multiple access control rules associated with file policies. You can configure file inspection for any Allow or Interactive Block access control rule, which permits you to match different file and malware inspection profiles against different types of traffic on your network before it reaches its final destination.

Before you begin

Complete the tasks up to this point in the malware protection workflow:

See How to Configure Malware Protection, on page 2214.

Procedure

Step 1 Review guidelines for file policies in access control policies. (These are different from the file rule and file policy guidelines that you looked at previously.)

Review File and Intrusion Inspection Order, on page 1730.

Step 2 Associate the file policy with an access control policy.

See Configuring an Access Control Rule to Perform Malware Protection, on page 2217

Step 3 Assign the access control policy to managed devices.

See Assigning Devices to an Access Control Policy, on page 1755.

What to do next

Continue with the next step in the malware protection workflow:

See How to Configure Malware Protection, on page 2214.

Configuring an Access Control Rule to Perform Malware Protection

Caution Enabling or disabling **Store files** in a **Detect Files** or **Block Files** rule, or adding the first or removing the last file rule that combines the **Malware Cloud Lookup** or **Block Malware** file rule action with an analysis option (**Spero Analysis or MSEXE**, **Dynamic Analysis**, or **Local Malware Analysis**) or a store files option (**Malware, Unknown**, **Clean**, or **Custom**), restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

Note Inline normalization is enabled automatically when a file policy is included in an access control rule. For more information, see https://www.cisco.com/go/snort3-inspectors.

Before you begin

- Adaptive profiling **must** be enabled (its default state) for access control rules to perform file control, including AMP.
- You must be an Admin, Access Admin, or Network Admin user to perform this task.

Procedure

Step 1 In the access control rule editor (from **Policies > Access Control)**, choose an **Action** of **Allow**, **Interactive Block**, or **Interactive Block with reset**.

Step 2 (Legacy UI only.) Click **Inspection**.

Step 3 Choose a **File Policy** to inspect traffic that matches the access control rule, or choose **None** to disable file inspection for matching traffic.

Step 4 (Optional) Disable logging of file or malware events for matching connections by clicking **Logging** and unchecking **Log Files**.

Note
Cisco recommends that you leave file and malware event logging enabled.

Step 5 Save the rule.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Related Topics

Create or Edit a File Policy, on page 2228
Snort Restart Scenarios, on page 252

Set Up Maintenance and Monitoring of Malware Protection

Ongoing maintenance is essential for protecting your network.

Before you begin

Configure your system to protect your network from malware.

See How to Configure Malware Protection, on page 2214 and referenced procedures.

Procedure

Step 1 Ensure that your system always has the most current and effective protection.

See Maintain Your System: File Types Eligible for Dynamic Analysis, on page 2227.

Step 2 Configure alerts for malware-related events and health monitoring.

See the Cisco Secure Firewall Management Center Administration Guide for information on *Configuring malware defense Alerting* and for information about the following modules:

- Local Malware Analysis

- Security Intelligence
- Threat Data Updates on Devices
- Intrusion and File Event Rate
- AMP for Firepower Status
- Secure Endpoint Status

What to do next

Review "What to do next items" in the malware protection workflow:

See How to Configure Malware Protection, on page 2214.

Cloud Connections for Malware Protection

Connections to public or private clouds are required in order to protect your network from malware.

AMP Clouds

The Advanced Malware Protection (AMP) cloud is a Cisco-hosted server that uses big data analytics and continuous analysis to provide intelligence that the system uses to detect and block malware on your network.

The AMP cloud provides dispositions for possible malware detected in network traffic by managed devices, as well as data updates for local malware analysis and file pre-classification.

If your organization has deployed Secure Endpoint and configured Firepower to import its data, the system imports this data from the AMP cloud, including scan records, malware detections, quarantines, and indications of compromise (IOC).

Cisco offers the following options for obtaining data from the Cisco cloud about known malware threats:

- **AMP public cloud**

  Your Secure Firewall Management Center communicates directly with the public Cisco cloud. There are three public AMP clouds, in the United States, Europe, and Asia.

- **An AMP private cloud**

  An AMP private cloud is deployed on your network and acts as a compressed, on-premises AMP cloud, as well as an anonymized proxy to connect to the public AMP cloud. For details, see Cisco AMP Private Cloud, on page 2221.

  If you integrate with Secure Endpoint, the AMP private cloud has some limitations. See Secure Endpoint and AMP Private Cloud, on page 2247.

Dynamic Analysis Cloud

- **Secure Malware Analytics Cloud**

  Public cloud that processes eligible files that you send for dynamic analysis, and provides threat scores and dynamic analysis reports. Firepower supports 200 samples/day for Secure Malware Analytics analysis.

- **On-premises Secure Malware Analytics Appliance**

  If your organization's security policy does not allow the system to send files outside of your network, you can deploy an on-premises appliance. This appliance does not contact the public Secure Malware Analytics Cloud.

  For more information, see Dynamic Analysis On-Premises Appliance (Cisco Secure Malware Analytics), on page 2225.

Configure Connections to AMP and Secure Malware Analytics Clouds

- AMP Cloud Connection Configurations, on page 2220
- Dynamic Analysis Connections, on page 2224

AMP Cloud Connection Configurations

The following topics describe AMP cloud connection configurations for different scenarios:

- Choose an AMP Cloud, on page 2221
- Connecting to an AMP Private Cloud, on page 2222
- Integrate Firepower and Secure Endpoint, on page 2247

The following topics are also relevant:

- Cisco AMP Private Cloud, on page 2221
- Requirements and Best Practices for AMP Cloud Connections, on page 2220
- Managing Connections to the AMP Cloud (Public or Private) , on page 2223

Requirements and Best Practices for AMP Cloud Connections

Requirements for AMP Cloud Connections

You must be an Admin user to set up the AMP cloud.

To ensure the Firewall Management Center can communicate with the AMP cloud, see *Security, Internet Access, and Communication Ports* in the Cisco Secure Firewall Management Center Administration Guide.

AMP and High Availability

Although they share file policies and related configurations, Firewall Management Centers in a high availability pair do not share AMP cloud connections, captured files, file events, and malware events. To ensure continuity of operations, and so that malware dispositions are the same on both Firewall Management Centers, both active and standby Firewall Management Centers must have access to the cloud.

These requirements apply to both public and private AMP clouds.

AMP Cloud Connections and Multitenancy

In a multidomain deployment, configure the malware defense connection at the Global level. Each Firewall Management Center can have only one malware defense connection.

Choose an AMP Cloud

By default, a connection to the United States (US) AMP public cloud is configured and enabled for your system. (This connection appears in the web interface as malware defense and sometimes AMP for Firepower.) You cannot delete or disable an malware defense cloud connection, but you can switch between different geographical AMP clouds, or configure an AMP private cloud connection.

Before you begin

- If you will use an AMP private cloud, see Connecting to an AMP Private Cloud, on page 2222 instead of this topic.
- Unless Firepower is integrated with Secure Endpoint, you can configure only one AMP cloud connection. This connection is labeled **AMP for Networks** or **AMP for Firepower**.
- If you have deployed Secure Endpoint and you want to add one or more AMP clouds to integrate that application with Firepower, see Integrate Firepower and Secure Endpoint, on page 2247.
- See Requirements and Best Practices for AMP Cloud Connections, on page 2220.

Procedure

Step 1 Choose **Integration** > **AMP** > **AMP Management**.

Step 2 Click pencil to edit the existing cloud connection.

Step 3 From the **Cloud Name** drop-down list, choose the regional cloud nearest to your Secure Firewall Management Center.

APJC is Asia/Pacific/Japan/China.

Step 4 Click **Save**.

What to do next

- If your deployment is a high-availability configuration, see Requirements and Best Practices for AMP Cloud Connections, on page 2220.
- (Optional) Change AMP Options, on page 2224.

Cisco AMP Private Cloud

The Firewall Management Center must connect to the AMP cloud for disposition queries for files detected in network traffic and receipt of retrospective malware events. This cloud can be public or private.

Your organization may have privacy or security concerns that make frequent or direct connections between your monitored network and the AMP cloud difficult or impossible. In these situations, you can set up a Cisco AMP Private Cloud, a proprietary Cisco product that acts as a compressed, on-premises version of the AMP cloud, as well as a secure mediator between your network and the AMP cloud. Connecting a Firewall Management Center to an AMP private cloud disables existing direct connections to the public AMP cloud.

All connections to the AMP cloud funnel through the AMP private cloud, which acts as an anonymized proxy to ensure the security and privacy of your monitored network. This includes disposition queries for files detected in network traffic, receiving of retrospective malware events, and so on. The AMP private cloud does not share any of your endpoint data over an external connection.

Note The AMP private cloud does **not** perform dynamic analysis, nor does it support anonymized retrieval of threat intelligence for other features that rely on Cisco Collective Security Intelligence (CSI), such as URL and Security Intelligence filtering.

For information about AMP private cloud (sometimes referred to as "AMPv"), see https://www.cisco.com/c/en/us/products/security/fireamp-private-cloud-virtual-appliance/index.html.

Connecting to an AMP Private Cloud

Before you begin

- Configure your Cisco AMP private cloud or clouds according to the directions in the documentation for that product. During configuration, note the private cloud host name. You will need this host name in order to configure the connection on the Firewall Management Center.
- Make sure the Firewall Management Center can communicate with the AMP private cloud, and confirm that the private cloud has internet access so it can communicate with the public AMP cloud. See the topics under *Security, Internet Access, and Communication Ports* in the Cisco Secure Firewall Management Center Administration Guide.
- Unless your deployment is integrated with Secure Endpoint, each Firewall Management Center can have only one AMP cloud connection. This connection is labeled **AMP for Networks** or **AMP for Firepower**.

  If you integrate with Secure Endpoint, you can configure multiple Secure Endpoint cloud connections.

Procedure

Step 1 Choose **Integration** > **AMP** > **AMP Management**.

Step 2 Click **Add AMP Cloud Connection**.

Step 3 From the **Cloud Name** drop-down list, choose **Private Cloud**.

Step 4 Enter a **Name**.

This information appears in malware events that are generated or transmitted by AMP private cloud.

Step 5 In the **Host** field, enter the private cloud host name that you configured when you set up the private cloud.

Step 6 Click **Browse** next to the **Certificate Upload Path** field to browse to the location of a valid TLS or SSL encryption certificate for the private cloud. For more information, see the AMP private cloud documentation.

Step 7 If you want to use this private cloud for both malware defense and Secure Endpoint, check the **Use for AMP for Firepower** check box.

If you configured a different private cloud to handle malware defense communications, you can clear this check box; if this is your only AMP private cloud connection, you cannot.

In a multidomain deployment, this check box appears only in the Global domain. Each Firewall Management Center can have only one malware defense connection.

Step 8 Click **Register**, confirm that you want to disable existing direct connections to the AMP cloud, and finally confirm that you want to continue to the AMP private cloud management console to complete registration.

Step 9 Log into the management console and complete the registration process. For further instructions, see the AMP private cloud documentation.

What to do next

In high availability deployments, configure AMP cloud connections on both management centers. These configurations are not synchronized.

Managing Connections to the AMP Cloud (Public or Private)

Use the Firewall Management Center to manage connections to public and private AMP clouds used for malware defense or Secure Endpoint or both.

You can delete a connection to a public or private AMP cloud if you no longer want to receive malware-related information from the cloud. Note that deregistering a connection using the Secure Endpoint or AMP private cloud management console does not remove the connection from the system. Deregistered connections display a failed state on the Secure Firewall Management Center web interface.

You can also temporarily disable a connection. When you reenable a cloud connection, the cloud resumes sending data to the system, including queued data from the disabled period.

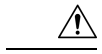

Caution For disabled connections, the public or private AMP cloud can store malware events, indications of compromise, and so on until you re-enable the connection. In rare cases—for example, with a very high event rate or a long-term disabled connection—the cloud may not be able to store all information generated while the connection is disabled.

In a multidomain deployment, the system displays connections created in the current domain, which you can manage. It also displays connections created in ancestor domains, which you cannot manage. To manage connections in a lower domain, switch to that domain. Each Firewall Management Center can have only one malware defense connection, which belongs to the Global domain.

Procedure

Step 1 Select **Integration** > **AMP** > **AMP Management**.

Step 2 Manage your AMP cloud connections:

- Delete — Click **Delete** ( ), then confirm your choice.
- Enable or Disable — Click the slider, then confirm your choice.

What to do next

In high availability deployments, configure AMP cloud connections on both management centers. These configurations are not synchronized.

Change AMP Options

Procedure

Step 1 Choose **Integration** > **Other Integrations** > **Cloud Services**.

Step 2 Select options:

Table 118: AMP for Networks Options

Option	Description
Enable Automatic Local Malware Detection Updates	The local malware detection engine statically analyzes and preclassifies files using signatures provided by Cisco. If you enable this option, the Firewall Management Center checks for signature updates once every 30 minutes.
Share URI from Malware Events with Cisco	The system can send information about the files detected in network traffic to the AMP cloud. This information includes URI information associated with detected files and their SHA-256 hash values. Although sharing is opt-in, transmitting this information to Cisco helps future efforts to identify and track malware.

Step 3 Click **Save**.

Dynamic Analysis Connections

Requirements for Dynamic Analysis

You must be an Admin, Access Admin, or Network Admin user, and be in the global domain, to use dynamic analysis.

With the appropriate license, the system automatically has access to the Secure Malware Analytics Cloud.

Dynamic analysis requires that managed devices have direct or proxied access to the Secure Malware Analytics Cloud or an on-premises Secure Malware Analytics Appliance on port 443.

See also Which Files Are Eligible for Dynamic Analysis?, on page 2239.

If you will connect to an on-premises Secure Malware Analytics Appliance, see also the prerequisites in Connect to an On-Premises Dynamic Analysis Appliance, on page 2225.

Viewing the Default Dynamic Analysis Connection

By default, the Secure Firewall Management Center can connect to the public Secure Malware Analytics Cloud for file submission and report retrieval. You can neither configure nor delete this connection.

Procedure

Step 1 Choose **Integration** > **AMP** > **Dynamic Analysis ConnectionsAMP** > **Dynamic Analysis Connections**.

Step 2 Click **Edit** ( ).

Note

For information about **Associate** ( ) on the **Integration** > **AMP** > **Dynamic Analysis Connections** page, see Enabling Access to Dynamic Analysis Results in the Public Cloud, on page 2226.

Dynamic Analysis On-Premises Appliance (Cisco Secure Malware Analytics)

If your organization has privacy or security concerns around submitting files to the public Secure Malware Analytics Cloud, you can deploy an on-premises Secure Malware Analytics Appliance. Like the public cloud, the on-premises appliance runs eligible files in a sandbox environment, and returns a threat score and dynamic analysis report to the system. However, the on-premises appliance does not communicate with the public cloud, or any other system external to your network.

For more information about on-premises Secure Malware Analytics Appliances, see https://www.cisco.com/c/en/us/products/security/threat-grid/index.html.

Connect to an On-Premises Dynamic Analysis Appliance

If you install an on-premises Secure Malware Analytics Appliance on your network, you can configure a dynamic analysis connection to submit files and retrieve reports from the appliance. When configuring the on-premises appliance dynamic analysis connection, you register the Secure Firewall Management Center to the on-premises appliance.

Before you begin

- Set up your on-premises Secure Malware Analytics Appliance.

  Documentation for this appliance is available from https://www.cisco.com/c/en/us/support/security/amp-threat-grid-appliances/tsd-products-support-series-home.html.

  For version requirements, see the *Cisco Firepower Compatibility Guide*.

- If your Secure Malware Analytics Appliance uses a self-signed public-key certificate, download the certificate from the Secure Malware Analytics Appliance; see the *Administrator's Guide* for your Secure Malware Analytics Appliance for information.

  If you use a certificate signed by a Certificate Authority (CA), the certificate must meet the following requirements:

  - The server key and signed certificate must be installed on the Secure Malware Analytics Appliance. Follow the upload instructions in the *Administrator's Guide* for your Secure Malware Analytics Appliance.
  - If there is a multi-level signing chain of CAs, all required intermediate certificates and the root certificate must be contained in a single file that will be uploaded to the Firewall Management Center.

- All certificates must be PEM-encoded.
- The file's newlines must be UNIX, not DOS.

- If you want to connect to the on-premises appliance using a proxy, configure the proxy; see *Modify Firewall Management Center Management Interfaces* in the Cisco Secure Firewall Management Center Administration Guide.
- Managed devices must have direct or proxied access to the Secure Malware Analytics Appliance on port 443.

Procedure

Step 1 Choose **Integration** > **AMP** > **Dynamic Analysis ConnectionsAMP** > **Dynamic Analysis Connections**.

Step 2 Click **Add New Connection**.

Step 3 Enter a **Name**.

Step 4 Enter a **Host**.

Step 5 Next to **Certificate Upload**, click **Browse** to upload the certificate for the on-premises appliance.

If the Secure Malware Analytics Appliance will present a self-signed certificate, upload the certificate you downloaded from that appliance.

If the Secure Malware Analytics Appliance will present a CA-signed certificate, upload the file containing the certificate signing chain.

Step 6 If you want to use a configured proxy to establish the connection, check the check box of **Use Proxy When Available**.

Step 7 Click **Register**.

Step 8 Click **Yes** to display the on-premises Secure Malware Analytics Appliance login page.

Step 9 Enter your username and password to the on-premises Secure Malware Analytics Appliance.

Step 10 Click **Sign in**.

Step 11 You have the following options:

- If you previously registered the Secure Firewall Management Center to the on-premises appliance, click **Return**.
- If you did not register the Secure Firewall Management Center, click **Activate**.

Enabling Access to Dynamic Analysis Results in the Public Cloud

Secure Malware Analytics offers more detailed reporting on analyzed files than is available in the Firewall Management Center. If your organization has a Secure Malware Analytics Cloud account, you can access the Secure Malware Analytics portal directly to view additional details about files sent for analysis from your managed devices. However, for privacy reasons, file analysis details are available only to the organization that submitted the files. Therefore, before you can view this information, you must associate your Firewall Management Center with the files submitted by its managed devices.

Before you begin

You must have a Secure Malware Analytics Cloud account, and have your account credentials ready.

Procedure

Step 1 Select **Integration** > **AMP** > **Dynamic Analysis ConnectionsAMP** > **Dynamic Analysis Connections**.

Step 2 Click **Associate** ( ) in the table row corresponding to the Secure Malware Analytics Cloud.

A Secure Malware Analytics portal window opens.

Step 3 Sign in to the Secure Malware Analytics Cloud.

Step 4 Click **Submit Query**.

Note
Do not change the default value in the **Devices** field.

If you have difficulties with this process, contact your Secure Malware Analytics representative at Cisco TAC.

It may take up to 24 hours for this change to take effect.

What to do next

After the association is activated, see *Viewing Dynamic Analysis Results in the Cisco Cloud* in the Cisco Secure Firewall Management Center Administration Guide .

Maintain Your System: File Types Eligible for Dynamic Analysis

The list of file types eligible for Dynamic Analysis is determined by the vulnerability database (VDB), which is updated periodically (but no more than once per day.) If you are an Admin user, you can update file types eligible for dynamic analysis.

The supported file types for automatic submission can vary in every release. The currently supported files for automatic submission are:

- PE Executable (32-bit Only)
- EXE
- DLL
- PDF
- MSOLE2 (Microsoft Object Linking and Embedding Compound File)
- DOCX, PPTX, XLSX

Note GZ and ZIP file types are not supported for automatic submission.

To ensure that your system has the current list:

Procedure

Step 1 Do one of the following:

- (Recommended) See *Vulnerability Database Update Automation* as discussed in the Cisco Secure Firewall Management Center Administration Guide
- Regularly check for new VDB updates, and *Manually Update the VDB* as discussed in the Cisco Secure Firewall Management Center Administration Guide when needed.

  If you choose this option, we recommend that you schedule regular reminders to do this.

Step 2 If your file policies specify individual file types instead of the **Dynamic Analysis Capable** file type category, update your file policies to use the newly supported file types.

Step 3 If the list of eligible file types changes, deploy to managed devices.

File Policies and File Rules

Create or Edit a File Policy

Before you begin

If you are configuring policies for malware protection, see all required procedures in Configure File Policies, on page 2216.

Procedure

Step 1 Select **Policies** > **Access Control heading** > **Malware & File** .

Step 2 Create a new policy, or edit an existing policy.

If you are editing an existing policy: If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Tip

To make a copy of an existing file policy, click **Copy** ( ), then type a unique name for the new policy in the dialog box that appears. You can then modify the copy.

Step 3 Add one or more rules to the file policy as described in Creating File Rules, on page 2241.

Step 4 Optionally, select Advanced and configure advanced options as described in Advanced and Archive File Inspection Options, on page 2229.

Step 5 Save the file policy.

What to do next

- If you are configuring policies for malware protection, see other required procedures in Configure File Policies, on page 2216.
- Otherwise:
  - Add the file policy to an access control rule as described in Add File Policies to Your Access Control Configuration, on page 2216.
  - Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Advanced and Archive File Inspection Options

The Advanced Settings in the file policy editor has the following general options:

- **First Time File Analysis**—Select this option to analyze first-seen files while AMP cloud disposition is pending. The file must match a rule configured to perform a malware cloud lookup and Spero, local malware, or dynamic analysis. If you deselect this option, files detected for the first time are marked with an Unknown disposition
- **Enable Custom Detection List**—Block files on the custom detection list.
- **Enable Clean List**—If enabled, this policy will allow files that are on the clean list.
- **If AMP Cloud disposition is Unknown, override disposition based upon threat score**—Select an option:
  - If you select **Disabled,** the system will not override the disposition provided by the AMP Cloud.
  - If you set a threshold threat score, files with an AMP cloud verdict of Unknown are considered malware if their Dynamic Analysis score is equal to or worse than the threshold.
  - If you select a lower threshold value, you increase the number of files treated as malware. Depending on the action selected in your file policy, this can result in an increase of blocked files.
  - For numeric threat score ranges, see *Threat Scores and Dynamic Analysis Summary Reports* in the Cisco Secure Firewall Management Center Administration Guide.

The Advanced Settings in the file policy editor has the following archive file inspection options:

- **Inspect Archives**—Enables inspection of the contents of archive files, for archive files as large as the **Maximum file size to store** advanced access control setting.
- **Block Encrypted Archives**—Blocks password-protected archives.
- **Block Uninspectable Archives**—Blocks archive files with contents that the system is unable to inspect for reasons other than encryption. This usually applies to corrupted files, or those that exceed your specified maximum archive depth.
- **Max Archive Depth**—Blocks nested archive files that exceed the specified depth. The top-level archive file is not considered in this count; depth begins at 1 with the first nested file .

Archive Files

Archive files are files that contain other files, such as `.zip` or `.rar` files.

If any individual file in an archive matches a file rule with a block action, the system blocks the entire archive, not just the individual file.

For details about options for archive file inspection, see Advanced and Archive File Inspection Options, on page 2229.

Archive Files That Can Be Inspected

- **File types**

  A complete list of inspectable archive file types appears in the Firewall Management Center web interface on the file rule configuration page. To view that page, see Creating File Rules, on page 2241.

  Contained files that can be inspected appears in the same page.

- **File size**

  You can inspect archive files as large as the **Maximum file size to store** file policy advanced access control setting.

- **Nested archives**

  Archive files can contain other archive files, which can in turn contain archive files. The level at which a file is nested is its *archive file depth*. Note that the top-level archive file is not included in the depth count; depth begins at 1 with the first nested file.

  The system can inspect up to three levels of nested files beneath the outermost archive file (level 0). You can configure your file policy to block archive files that exceed that depth (or a lower maximum depth that you specify).

  If you choose not to block files that exceed the maximum archive file depth of 3, when archive files that contain some extractable contents and some contents nested at a depth of 3 or greater appear in monitored traffic, the system examines and reports data only for the files it was able to inspect.

  All features applicable to uncompressed files (such as dynamic analysis and file storage) are available for nested files inside archive files.

- **Encrypted files**

  You can configure the system to block archives whose contents are encrypted or otherwise cannot be inspected.

- **Archives that are not inspected**

  If traffic that contains an archive file is on a Security Intelligence Block list or Do Not Block list, or if the top-level archive file's SHA-256 value is on the custom detection list, the system does not inspect the contents of the archive file.

  If a nested file is blocked, the entire archive is blocked; however, if a nested file is allowed, the archive is not automatically passed (depending on any other nested files and characteristics).

  .Exe files inside some .rar archives cannot be detected, including possibly rar5.

Archive File Dispositions

Archive file dispositions are based on the dispositions assigned to the files inside the archive. **All** archives that contain identified malware files receive a disposition of `Malware`. Archives without identified malware files receive a disposition of `Unknown` if they contain any unknown files, and a disposition of `Clean` if they contain only clean files.

Table 119: Archive File Disposition by Contents

Archive File Disposition	Number of Unknown Files	Number of Clean Files	Number of Malware Files
`Unknown`	1 or more	Any	0
`Clean`	0	1 or more	0
`Malware`	Any	Any	1 or more

Archive files, like other files, may have dispositions of `Custom Detection` or `Unavailable` if the conditions for those dispositions apply.

Viewing Archive Contents and Details

If your file policy is configured to inspect archive file contents, you can use the context menu in a table on pages under the Analysis > Files menu, and the network file trajectory viewer to view information about the files inside an archive when the archive file appears in a file event, malware event, or as a captured file.

All file contents of the archive are listed in table form, with a short summary of their relevant information: name, SHA-256 hash value, type, category, and archive depth. A network file trajectory icon appears by each file, which you can click to view further information about that specific file.

Override File Disposition Using Custom Lists

If a file has a disposition in the AMP cloud that you know to be incorrect, you can add the file's SHA-256 value to a file list that overrides the disposition from the cloud:

- To treat a file as if the AMP cloud assigned a clean disposition, add the file to the *clean list*.
- To treat a file as if the AMP cloud assigned a malware disposition, add the file to the *custom detection list.*

On subsequent detection, the device either allows or blocks the file without reevaluating the file's disposition. You can use the clean list or custom detection list per file policy.

Note To calculate a file's SHA-256 value, you must configure a rule in the file policy to either perform a malware cloud lookup or block malware on matching files.

For complete information about using file lists in Firepower, see File List, on page 1392.

Alternatively, if applicable, use Centralized File Lists from Secure Endpoint, on page 2231.

Centralized File Lists from Secure Endpoint

If your organization has deployed Secure Endpoint, Firepower can use Block and Allow lists created in Secure Endpoint when it queries the AMP cloud for file dispositions.

Requirements:

- Your organization must be using the AMP public cloud.
- Your organization has deployed Secure Endpoint.

- You have registered your system to Secure Endpoint using the procedure in Integrate Firepower and Secure Endpoint, on page 2247.

To create and deploy these lists, see the documentation or online help for Secure Endpoint.

Note File lists created in Firepower override file lists created in Secure Endpoint.

Managing File Policies

The File Policies page displays a list of existing file policies along with their last-modified dates. You can use this page to manage your file policies.

Note The system checks for updates to the list of file types eligible for dynamic analysis (no more than once a day). If the list of eligible file types changes, this constitutes a change in the file policy; any access control policy using the file policy is marked out-of-date if deployed to any devices. You must deploy policies before the updated file policy can take effect on the device. See Maintain Your System: File Types Eligible for Dynamic Analysis, on page 2227.

Procedure

Step 1 Select **Policies** > **Access Control heading** > **Malware & File** .

Step 2 Manage your file policies:

- Compare—Click **Compare Policies**; see Comparing policies.
- Create — To create a file policy, click **New File Policy** and proceed as described in Create or Edit a File Policy, on page 2228.
- Copy — To copy a file policy, click **Copy** ( ).

  If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Delete — If you want to delete a file policy, click **Delete** ( ), then click **Yes** and **OK** as prompted.

  If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Deploy—Choose **Deploy** > **Deployment**; see Deploy Configuration Changes, on page 259.
- Edit — If you want to modify an existing file policy, click **Edit** ( ).
- Report—Click **Report**( ); see Generate Current Policy Reports, on page 278.

File Rules

A file policy, like its parent access control policy, contains rules that determine how the system handles files that match the conditions of each rule. You can configure separate file rules to take different actions for different file types, application protocols, or directions of transfer.

For example, when a file matches a rule, the rule can:

- allow or block files based on simple file type matching
- block files based on disposition (whether or not evaluation indicates that it is malicious)
- store files to the device (For information, see Captured Files and File Storage, on page 2239)
- submit stored (captured) files for local malware, Spero, or dynamic analysis

In addition, the file policy can:

- automatically treat a file as if it is clean or malware based on entries in the clean list or custom detection list
- treat a file as if it is malware if the file's threat score exceeds a configurable threshold
- inspect the contents of archive files (such as `.zip` or `.rar`)
- block archive files whose contents are encrypted, nested beyond a specified maximum archive depth, or otherwise uninspectable

File Rule Components

Table 120: File Rule Components

File Rule Component	Description
application protocol	The system can detect and inspect files transmitted via FTP, HTTP, SMTP, IMAP, POP3, and NetBIOS-ssn (SMB). **Any**, the default, detects files in HTTP, SMTP, IMAP, POP3, FTP, and NetBIOS-ssn (SMB) traffic. To improve performance, you can restrict file detection to only one of those application protocols on a per-file rule basis.
direction of transfer	You can inspect incoming FTP, HTTP, IMAP, POP3, and NetBIOS-ssn (SMB) traffic for downloaded files; you can inspect outgoing FTP, HTTP, SMTP, and NetBIOS-ssn (SMB) traffic for uploaded files. **Tip** Use **Any** to detect files over multiple application protocols, regardless of whether users are sending or receiving.

File Rule Component	Description
file categories and types	The system can detect various types of files. These file types are grouped into basic categories, including multimedia (swf, mp3), executables (exe, torrent), and PDFs. You can configure file rules that detect individual file types, or on entire categories of file types. For example, you could block all multimedia files, or just ShockWave Flash (swf) files. Or, you could configure the system to alert you when a user downloads a BitTorrent (torrent) file. Note that executables include file types that can run macros and scripts, since these can contain malware. For a list of file types the system can inspect, select **Policies** > **Access Control heading** > **Malware & File** , create a temporary new file policy, then click **Add Rule**. Select a file type category and the file types that the system can inspect appear in the **File Types** list. **Note** Frequently triggered file rules can affect system performance. For example, detecting multimedia files in HTTP traffic (YouTube, for example, transmits significant Flash content) could generate an overwhelming number of events.
file rule action	A file rule's action determines how the system handles traffic that matches the conditions of the rule. Depending on the selected action, you can configure whether the system stores the file or performs Spero, local malware, or dynamic analysis on a file. If you select a Block action, you can also configure whether the system also resets the blocked connection. For descriptions of these actions and options, see File Rule Actions, on page 2234. File rules are evaluated in rule-action, not numerical, order. For details, see File Rule Actions: Evaluation Order, on page 2241.

File Rule Actions

File rules give you granular control over which file types you want to log, block, or scan for malware. Each file rule has an associated action that determines how the system handles traffic that matches the conditions of the rule. To be effective, a file policy must contain one or more rules. You can use separate rules within a file policy to take different actions for different file types, application protocols, or directions of transfer.

File Rule Actions

- *Detect Files* rules allow you to log the detection of specific file types to the database, while still allowing their transmission.
- *Block Files* rules allow you to block specific file types. You can configure options to reset the connection when a file transfer is blocked, and store captured files to the managed device.
- *Malware Cloud Lookup* rules allow you to obtain and log the disposition of files traversing your network, while still allowing their transmission.

- *Block Malware* rules allow you to calculate the SHA-256 hash value of specific file types, query the AMP cloud to determine if files traversing your network contain malware, then block files that represent threats.

File Rule Action Options

Depending on the action you select, you have different options:

File Rule Action Option	Block Files capable?	Block Malware capable?	Detect Files capable?	Malware Cloud Lookup capable?
Spero Analysis* for MSEXE	no	yes, you can submit executable files	no	yes, you can submit executable files
Dynamic Analysis*	no	yes, you can submit executable files with Unknown file dispositions	no	yes, you can submit executable files with Unknown file dispositions
Capacity Handling	no	yes	no	yes
Local Malware Analysis*	no	yes	no	yes
Reset Connection	yes (recommended)	yes (recommended)	no	no
Store files	yes, you can store all matching file types	yes, you can store file types matching the file dispositions you select	yes, you can store all matching file types	yes, you can store file types matching the file dispositions you select

* For complete information about these options, see Malware Protection Options (in File Rule Actions), on page 2235 and its subtopics.

Caution Enabling or disabling **Store files** in a **Detect Files** or **Block Files** rule, or adding the first or removing the last file rule that combines the **Malware Cloud Lookup** or **Block Malware** file rule action with an analysis option (**Spero Analysis or MSEXE**, **Dynamic Analysis**, or **Local Malware Analysis**) or a store files option (**Malware, Unknown**, **Clean**, or **Custom**), restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

Malware Protection Options (in File Rule Actions)

The system applies several methods of file inspection and analysis to determine whether a file contains malware.

Depending on the options you enable in a file rule, the system inspects files using the following tools, in order:

1. Spero Analysis, on page 2237 and AMP Cloud Lookup, on page 2238
2. Local Malware Analysis, on page 2238

3. Dynamic Analysis, on page 2238

For a comparison of these tools, see Comparison of Malware Protection Options, on page 2236.

(You can also, if you choose, block all files based on their file type. For more information, see Block All Files by Type, on page 2241.)

See also information about Cisco's Secure Endpoint product at (Optional) Malware Protection with Secure Endpoint, on page 2245 and subtopics.

Comparison of Malware Protection Options

The following table details the benefits and drawbacks of each type of file analysis, as well as the way each malware protection method determines a file's disposition.

Analysis Type	Benefit	Limitations	Malware Identification
Spero analysis	Structural analysis of executable files, submits Spero signature to the AMP Cloud for analysis	Less thorough than local malware analysis or dynamic analysis, only for executable files	Disposition changes from Unknown to Malware only on positive identification of malware.
Local malware analysis	Consumes fewer resources than dynamic analysis, and returns results more quickly, especially if the detected malware is common	Less thorough results than dynamic analysis	Disposition changes from Unknown to Malware only on positive identification of malware.
Dynamic analysis	Thorough analysis of unknown files using Secure Malware Analytics	Eligible files are uploaded to the public cloud or an on-premises appliance. It takes some time to complete analysis	Threat score determines maliciousness of a file. Disposition can be based on the threat score threshold configured in the file policy.
Spero analysis and local malware analysis	Consumes fewer resources than configuring local malware analysis and dynamic analysis, while still using AMP cloud resources to identify malware	Less thorough than dynamic analysis, Spero analysis only for executable files	Disposition changes from Unknown to Malware only on positive identification of malware.

Analysis Type	Benefit	Limitations	Malware Identification
Spero analysis and dynamic analysis	Uses full capabilities of AMP cloud in submitting files and Spero signatures	Results obtained less quickly than if using local malware analysis	Threat score changes based on dynamic analysis results for files preclassified as possible malware. Disposition changes based on configured threat score threshold in the file policy, and from Unknown to Malware if the Spero analysis identifies malware.
Local malware analysis and dynamic analysis	Thorough results in using both types of file analysis	Consumes more resources than either alone	Threat score changes based on dynamic analysis results for files preclassified as possible malware. Disposition changes from Unknown to Malware if local malware analysis identifies malware, or based on configured threat score threshold in the file policy.
Spero analysis, local malware analysis and dynamic analysis	Most thorough results	Consumes most resources in running all three types of file analysis	Threat score changes based on dynamic analysis results for files preclassified as possible malware. Disposition changes from Unknown to Malware if Spero analysis or local malware analysis identifies malware, or based on configured threat score threshold in the file policy.
(Block transmission of all files of a specified file type)	Does not require a Malware Defense license (This option is not technically a malware protection option.)	Legitimate files will also be blocked	(No analysis is performed.)

Note Preclassification does not itself determine a file's disposition; it is merely one of the factors that determine whether a file is eligible for Dynamic Analysis.

Spero Analysis

Spero analysis examines structural characteristics such as metadata and header information in executable files. After generating a Spero signature based on this information, if the file is an eligible executable file, the device submits it to the Spero heuristic engine in the AMP cloud. Based on the Spero signature, the Spero engine determines whether the file is malware. You can also configure rules to submit files for Spero analysis without also submitting them to the AMP cloud.

Note that you cannot manually submit files for Spero analysis.

AMP Cloud Lookup

For files that are eligible for assessment using Advanced Malware Protection, the Firewall Management Center performs a *malware cloud lookup*, querying the AMP cloud for the file's disposition based on its SHA-256 hash value.

To improve performance, the system caches dispositions returned by the cloud and uses the cached disposition for known files rather than querying the AMP cloud. For more information about this cache, see Cached Disposition Longevity, on page 2238.

Local Malware Analysis

Local malware analysis allows a managed device to locally inspect executables, PDFs, office documents, and other types of files for the most common types of malware, using a detection rule set provided by the Talos Intelligence Group. Because local analysis does not query the AMP cloud, and does not run the file, local malware analysis saves time and system resources.

If the system identifies malware through local malware analysis, it updates the existing file disposition from Unknown to Malware. The system then generates a new malware event. If the system does not identify malware, it does not update the file disposition from Unknown to Clean. After the system runs local malware analysis, it caches file information such as SHA-256 hash value, timestamp, and disposition, so that if detected again within a certain period of time, the system can identify malware without additional analysis. For more information about the cache, see Cached Disposition Longevity, on page 2238.

Local malware analysis does not require establishing communications with the Secure Malware Analytics Cloud. However, you must configure communications with the cloud to submit files for dynamic analysis, and to download updates to the local malware analysis ruleset.

Cached Disposition Longevity

Dispositions returned from an AMP cloud query, associated threat scores, and dispositions assigned by local malware analysis, have a time-to-live (TTL) value. After a disposition has been held for the duration specified in the TTL value without update, the system purges the cached information. Dispositions and associated threat scores have the following TTL values:

- Clean — 4 hours
- Unknown — 1 hour
- Malware — 1 hour

If a query against the cache identifies a cached disposition that timed out, the system re-queries the local malware analysis database and the AMP cloud for a new disposition.

Dynamic Analysis

You can configure your file policy to automatically submit files for dynamic analysis using Secure Malware Analytics (formerly Threat Grid), Cisco’s file analysis and threat intelligence platform.

Devices submit eligible files to Secure Malware Analytics (either the public cloud or to an on-premises appliance, whichever you have specified) regardless of whether the device stores the file.

Secure Malware Analytics runs the file in a sandbox environment, analyzes the file's behavior to determine whether the file is malicious, and returns a threat score that indicates the likelihood that a file contains malware. From the threat score, you can view a dynamic analysis summary report with the reasons for the assigned

threat score. You can also look in Secure Malware Analytics to view detailed reports for files that your organization submitted, as well as scrubbed reports with limited data for files that your organization did not submit.

For more information about Cisco Secure Malware Analytics, see https://www.cisco.com/c/en/us/products/security/threat-grid/index.html

To configure your system to perform dynamic analysis, see the topics under Dynamic Analysis Connections, on page 2224.

Which Files Are Eligible for Dynamic Analysis?

A file's eligibility for dynamic analysis depends on:

- the file type
- the file size
- the file rule's action

Additionally:

- The system submits only files that match the file rules you configure.
- The file must have a malware cloud lookup disposition of Unknown or Unavailable at the time the file is sent for analysis.
- The system must preclassify the file as potential malware.

Dynamic Analysis and Capacity Handling

Capacity handling allows you to temporarily store files that are otherwise eligible for dynamic analysis if the system is temporarily unable to submit files to the cloud, either because the device cannot communicate with the cloud or because the maximum number of submissions has been reached. The system submits the stored files when the hindering condition has passed.

Some devices can store files on the device hard drive or in a malware storage pack. See also Malware Storage Pack, on page 2240.

Captured Files and File Storage

The file storage feature allows you to capture selected files detected in traffic, and automatically store a copy of the file temporarily to a device’s hard drive, or, if installed, to the malware storage pack.

After your device captures the files, you can:

- Store captured files on the device’s hard drive for later analysis.
- Download the stored file to a local computer for further manual analysis or archival purposes.
- Manually submit eligible captured files for AMP cloud lookup or dynamic analysis.

Note that once a device stores a file, it will not re-capture it if the file is detected in the future and the device still has that file stored.

Note When a file is detected for the first time on your network, you can generate a file event that represents the file's detection. However, if your file rule performs a malware cloud lookup, the system requires additional time to query the AMP cloud and return a disposition. Due to this delay, the system cannot store this file until the second time it is seen on your network, and the system can immediately determine the file's disposition.

Whether the system captures or stores a file, you can:

- Review information about the captured file from Analysis > Files > Captured Files, including whether the file was stored or submitted for dynamic analysis, file disposition, and threat score, allowing you to quickly review possible malware threats detected on your network.
- View the file's trajectory to determine how it traversed your network and which hosts have a copy.
- Add the file to the clean list or custom detection list to always treat the file as if it had a clean or malware disposition on future detection.

You configure file rules in a file policy to capture and store files of a specific type, or with a particular file disposition, if available. After you associate the file policy with an access control policy and deploy it to your devices, matching files in traffic are captured and stored. You can also limit the minimum and maximum file sizes to store.

Stored files are not included in system backups.

You can view captured file information under Analysis > Files > Captured Files, and download a copy for offline analysis.

Malware Storage Pack

Based on your file policy configuration, your device may store a substantial amount of file data to the hard drive. You can install a malware storage pack in the device; the system stores files to the malware storage pack, allowing more room on the primary hard drive to store events and configuration files. The system periodically deletes older files. If the device's primary hard drive does not have enough available space, and does not have an installed malware storage pack, you cannot store files.

Caution Do not attempt to install a hard drive that was not supplied by Cisco in your device. Installing an unsupported hard drive may damage the device. Malware storage pack kits are available for purchase **only** from Cisco. Contact Support if you require assistance with the malware storage pack.

Without a malware storage pack installed, when you configure a device to store files, it allocates a set portion of the primary hard drive's space to captured file storage. If you configure capacity handling to temporarily store files for dynamic analysis, the system uses the same hard drive allocation to store these files until it can resubmit them to the cloud.

When you install a malware storage pack in a device and configure file storage or capacity handling, the device allocates the entire malware storage pack for storing these files. The device cannot store any other information on the malware storage pack.

When the allocated space for captured file storage fills to capacity, the system deletes the oldest stored files until the allocated space reaches a system-defined threshold. Based on the number of files stored, you may see a substantial drop in disk usage after the system deletes files.

If a device has already stored files when you install a malware storage pack, the next time you restart the device, any captured files or capacity handling files stored on the primary hard drive are moved to the malware storage pack. Any future files the device stores are stored to the malware storage pack.

For more information on using MSP on the Firepower devices, see the Firepower Hardware Installation Guide for your device.

Block All Files by Type

If your organization wants to block not only the transmission of malware files, but all files of a specific type, regardless of whether the files contain malware, you can do so.

File control is supported for all file types where the system can detect malware, plus many additional file types. These file types are grouped into basic categories, such as multimedia (swf, mp3), executables (exe, torrent), and PDFs.

Blocking all files based on their type is not technically a malware protection feature; it does not require a Malware Defense license and does not query the AMP cloud.

File Rule Actions: Evaluation Order

A file policy will likely contain multiple rules with different actions for different situations. If more than one rule can apply to a particular situation, the evaluation order described in this topic applies. In general, simple blocking takes precedence over malware inspection and blocking, which takes precedence over simple detection and logging.

The order of precedence of file-rule actions is:

- *Block Files*
- *Block Malware*
- *Malware Cloud Lookup*
- *Detect Files*

If configured, TID also impacts action prioritization. For more information, see Threat Intelligence Director-Firewall Management Center Action Prioritization, on page 2757.

Creating File Rules

Caution Enabling or disabling **Store files** in a **Detect Files** or **Block Files** rule, or adding the first or removing the last file rule that combines the **Malware Cloud Lookup** or **Block Malware** file rule action with an analysis option (**Spero Analysis or MSEXE**, **Dynamic Analysis**, or **Local Malware Analysis**) or a store files option (**Malware, Unknown**, **Clean**, or **Custom**), restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

Before you begin

If you are configuring rules for malware protection, see Configure File Policies, on page 2216.

Procedure

Step 1 Select **Policies** > **Access Control heading** > **Malware & File** .

Step 2 Click the edit icon to modify an existing file policy.

Step 3 In the file policy editor, click **Add Rule**.

Step 4 Select an **Application Protocol** and **Direction of Transfer** as described in File Rule Components, on page 2233.

Step 5 Select one or more **File Types**.

The file types you see depend on the selected application protocol, direction of transfer, and action.

You can filter the list of file types in the following ways:

- Select one or more **File Type Categories**, then click **All types in selected Categories.**
- Search for a file type by its name or description. For example, type **`Windows`** in the **Search name and description** field to display a list of Microsoft Windows-specific files.

Tip
Hover your pointer over a file type to view its description.

Step 6 Select a file rule **Action** as described in File Rule Actions, on page 2234, with consideration for File Rule Actions: Evaluation Order, on page 2241.

The actions available to you depend on the licenses you have installed. See License Requirements for File and Malware Policies, on page 2211.

Step 7 Depending on the action you selected, configure options:

- reset the connection after blocking the file
- store files that match the rule
- enable Spero analysis*
- enable local malware analysis*
- enable dynamic analysis* and capacity handling

* For information about these options, see File Rule Actions, on page 2234 and Malware Protection Options (in File Rule Actions), on page 2235 and its subtopics.

Step 8 Click **Add.**

Step 9 Click **Save** to save the policy.

What to do next

- If you are configuring policies for malware protection, return to Configure File Policies, on page 2216.
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Access Control Rule Logging for Malware Protection

When the system detects a prohibited file (including malware) according to the settings in the file policy, it automatically logs an event to the Secure Firewall Management Center database. If you do not want to log file or malware events, you can disable this logging on a per-access-control-rule basis.

The system also logs the end of the associated connection to the Secure Firewall Management Center database, regardless of the logging configuration of the invoking access control rule.

Retrospective Disposition Changes

File dispositions can change. For example, as new information is discovered, the AMP cloud can determine that a file that was previously thought to be clean is now identified as malware, or the reverse—that a malware-identified file is actually clean. When the disposition changes for a file you queried in the past week, the AMP cloud notifies the system so it can automatically take action the next time it detects that file being transmitted. A changed disposition is called a *retrospective* disposition.

File and Malware Inspection Performance and Storage Options

Increasing the file sizes can affect the performance of the system.

Table 121: Advanced Access Control File and malware defense Options

Field	Description	Guidelines and Restrictions
Limit the number of bytes inspected when doing file type detection	Specifies the number of bytes inspected when performing file type detection.	0 - 4294967295 (4GB) 0 removes the restriction. The default value is the maximum segment size of a TCP packet (1460 bytes). In most cases, the system can identify common file types using the first packet. To detect ISO files, enter a value greater than 36870.
Allow file if cloud lookup for Block Malware takes longer than (seconds)	Specifies how long the system will hold the last byte of a file that matches a **Block Malware** rule and that does not have a cached disposition, while malware cloud lookup occurs. If the time elapses without the system obtaining a disposition, the file passes. Dispositions of Unavailable are not cached.	0 - 30 seconds Do *not* set this option to 0 without contacting Support. Cisco recommends that you use the default value to avoid blocking traffic because of connection failures.
Do not calculate SHA-256 hash values for files larger than (in bytes)	Prevents the system from storing files larger than a certain size, performing a malware cloud lookup on the files, or blocking the files if added to the custom detection list.	0 - 4294967295 (4GB) 0 removes the restriction. This value must be greater than or equal to **Maximum file size to store (bytes)** and **Maximum file size for dynamic analysis testing (bytes)**.

Field	Description	Guidelines and Restrictions

Minimum file size for advanced file inspection and storage (bytes)	These settings specify:
	• The file size that the system can inspect using the following detectors:
	  • Spero analysis
	  • Sandboxing and preclassification
	  • Local malware analysis/ClamAV
	  • Archive inspection
	• The file size that the system can store using a file rule.
		0 - 10485760 (10MB)
		0 disables file storage.
		Must be less than or equal to **Maximum file size to store (bytes)** and **Do not calculate SHA-256 hash values for files larger than (in bytes)**.

Maximum file size for advanced file inspection and storage (bytes)		0 - 10485760 (10MB) 0 disables file storage. Must be greater than or equal to **Minimum file size to store (bytes)**, and less than or equal to **Do not calculate SHA-256 hash values for files larger than (in bytes)**.
Minimum file size for dynamic analysis testing (bytes)	Specifies the minimum file size the system can submit to the AMP cloud for dynamic analysis.	0 -10485760 (10MB) Must be less than or equal to **Maximum file size for dynamic analysis testing (bytes)** and **Do not calculate SHA-256 hash values for files larger than (in bytes)**. The file size for dynamic analysis must be within the limits defined by the minimum and maximum settings for file analysis. The system checks the AMP cloud for updates to the minimum file size you can submit (no more than once a day). If the new minimum size is larger than your current value, your current value is updated to the new minimum, and your policy is marked out-of-date.
Maximum file size for dynamic analysis testing (bytes)	Specifies the maximum file size the system can submit to the AMP cloud for dynamic analysis.	0 -10485760 (10MB) Must be greater than or equal to **Minimum file size for dynamic analysis testing (bytes)**, and less than or equal to **Do not calculate SHA-256 hash values for files larger than (in bytes)**. The file size for dynamic analysis must be within the limits defined by the minimum and maximum settings for file analysis. The system checks the AMP cloud for updates to the maximum file size you can submit (no more than once a day). If the new maximum size is smaller than your current value, your current value is updated to the new maximum, and your policy is marked out-of-date.

Tuning File and Malware Inspection Performance and Storage

You must be an Admin, Access Admin, or Network Admin user to perform this task.

Procedure

Step 1 In the access control policy editor, click **Advanced Settings**.

Step 2 Click **Edit** (✎) next to **Files and Malware Settings**.

If **View** (👁) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Set any of the options described in File and Malware Inspection Performance and Storage Options, on page 2243.

Step 4 Click **OK**.

Step 5 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

(Optional) Malware Protection with Secure Endpoint

Cisco's Secure Endpoint is a separate malware-protection product that can supplement malware protection provided by the system and be integrated with your Firepower deployment.

Secure Endpoint is Cisco’s enterprise-class Advanced Malware Protection solution that runs as a lightweight connector on individual users' *endpoints* (computers and mobile devices) to discover, understand, and block advanced malware outbreaks, advanced persistent threats, and targeted attacks.

Benefits of Secure Endpoint include:

- Configure custom malware detection policies and profiles for your entire organization, as well as perform flash and full scans on all your users’ files
- Perform malware analysis, including view heat maps, detailed file information, network file trajectory, and threat root causes
- Configure multiple aspects of outbreak control, including automatic quarantines, application blocking to stop non-quarantined executables from running, and exclusion lists
- Create custom protections, block execution of certain applications based on group policy, and create custom Allowed Applications lists
- Use the Secure Endpoint management console to help you mitigate the effect of malware. The management console provides a robust, flexible web interface where you control all aspects of your Secure Endpoint deployment and manage all phases of an outbreak.

For detailed information about Secure Endpoint, see:

- https://www.cisco.com/c/en/us/products/security/amp-for-endpoints/index.html.
- Online help in the Secure Endpoint management console.

- Secure Endpoint documentation available from: http://docs.amp.cisco.com.

Comparison of Malware Protection: Firepower vs. Secure Endpoint

Table 122: Advanced Malware Protection Differences by Detecting Product

Feature	Firepower Malware Protection (malware defense)	Secure Endpoint
File type detection and blocking method (file control)	In network traffic, using access control and file policies	Not supported
Malware detection and blocking method	In network traffic, using access control and file policies	On individual endpoints (end-user computers and mobile devices), using a connector that communicates with the AMP cloud
Network traffic inspected	Traffic passing through a managed device	None; connectors installed on endpoints directly inspect files
Malware intelligence data source	AMP cloud (public or private)	AMP cloud (public or private)
Malware detection robustness	Limited file types	All file types
Malware analysis choices	Firewall Management Center-based, plus analysis in the AMP cloud	Firewall Management Center-based, plus additional options on the Secure Endpoint management console
Malware mitigation	Malware blocking in network traffic, Firewall Management Center-initiated remediations	Secure Endpoint-based quarantine and outbreak control options, Firewall Management Center-initiated remediations
Events generated	File events, captured files, malware events, and retrospective malware events	Malware events
Information in malware events	Basic malware event information, plus connection data (IP address, port, and application protocol)	In-depth malware event information; no connection data
Network file trajectory	Firewall Management Center-based	Firewall Management Center and the Secure Endpoint management console each have a network file trajectory. Both are useful.
Required licenses or subscriptions	Licenses required to perform file control and malware defense	Secure Endpoint subscription. No license is required to bring Secure Endpoint data into Firewall Management Center.

About Integrating Firepower with Secure Endpoint

If your organization has deployed Secure Endpoint, you can optionally integrate that product with your Firepower deployment.

Integration with Secure Endpoint does not require a dedicated Firepower license.

Benefits of Integrating Firepower and Secure Endpoint

Integrating your Secure Endpoint deployment with your system offers the following benefits:

- Centralized Blocked Applications and Allowed Applications lists configured in Secure Endpoint can determine verdicts for file SHAs sent from Firepower to the AMP cloud for disposition.

  See Centralized File Lists from Secure Endpoint, on page 2231.

- The system can import malware events detected by Secure Endpoint into Secure Firewall Management Center so you can manage these events along with malware events generated by the system. Imported data for these events includes scans, malware detections, quarantines, blocked executions, and cloud recalls, as well as indications of compromise (IOCs) that Firewall Management Center displays for hosts that it monitors.

  For more information, see *Malware Event Analysis with Secure Endpoint* in the Cisco Secure Firewall Management Center Administration Guide .

- You can view file trajectory and other details in the Secure Endpoint console.

  For details, see *Work with Event Data in the Secure Endpoint Console* in the Cisco Secure Firewall Management Center Administration Guide .

Important If you use a Cisco AMP Private Cloud, see limitations at Secure Endpoint and AMP Private Cloud, on page 2247.

Secure Endpoint and AMP Private Cloud

If you configure a Cisco AMP private cloud to collect Secure Endpoint data on your network, all Secure Endpoint connectors send data to the private cloud, which forwards that data to the Secure Firewall Management Center. The private cloud does not share any of your endpoint data over an external connection.

If your organization has deployed an AMP private cloud, all connections to the AMP cloud funnel through the private cloud, which acts as an anonymized proxy to ensure the security and privacy of your monitored network. This includes importing Secure Endpoint data. The private cloud does not share any of your endpoint data over an external connection.

The following integration features are not available if you use an AMP private cloud:

- Use of Blocked Applications and Allowed Applications lists configured in Secure Endpoint. (These lists are used to block or allow files.)
- Visibility in Secure Endpoint of malware events generated from Firepower.

You can configure multiple private clouds to support the capacity you require.

Integrate Firepower and Secure Endpoint

If your organization has deployed Cisco's Secure Endpoint product, you can integrate that application with Firepower to achieve the benefits described in Benefits of Integrating Firepower and Secure Endpoint, on page 2247.

When you integrate with Secure Endpoint, you must configure the Secure Endpoint connection even if you already have malware defense (AMP for Firepower) connections configured. You can configure multiple Secure Endpoint cloud connections.

Note Secure Endpoint connections that have not registered successfully do not affect malware defense.

Before you begin

- You must be an Admin user to perform this task.
- If your deployment uses Cisco AMP Private Cloud, see limitations at Secure Endpoint and AMP Private Cloud, on page 2247.
- Secure Endpoint must be set up and working properly on your network.
- The Firewall Management Center must have internet access.
- Make sure your Firewall Management Center and Secure Endpoint can communicate with each other. See the topics under *Security, Internet Access, and Communication Ports* in the Cisco Secure Firewall Management Center Administration Guide.
- If you are connecting to the AMP cloud after either reimaging or restoring your Firewall Management Center from backup, use the Secure Endpoint management console to remove the previous connection.
- You will need your Secure Endpoint credentials to log in to the Secure Endpoint console during this procedure.

Procedure

Step 1 Choose **Integration** > **AMP** > **AMP Management**.

Step 2 Click **Add AMP Cloud Connection**.

Step 3 From the **Cloud Name** drop-down list, choose the cloud you want to use.

- The AMP cloud closest to the geographical location of your Firewall Management Center.

  APJC is Asia/Pacific/Japan/China.

- For AMP private cloud (AMPv), choose **Private Cloud** and proceed as described in Cisco AMP Private Cloud, on page 2221.

Step 4 If you want to use this cloud for both malware defense and Secure Endpoint, select the **Use for AMP for Firepower** check box.

If you configured a different cloud to handle malware defense (AMP for Firepower) communications, you can clear this check box; if this is your only AMP cloud connection, you cannot.

Step 5 Click **Register**.

A **Spinning state** ( ) icon indicates that a connection is pending, for example, after you configure a connection on the Firewall Management Center, but before you authorize it using the Secure Endpoint management

console. A **Denied** ( ) icon indicates that the cloud denied the connection or the connection failed for another reason.

Step 6 Confirm that you want to continue to the Secure Endpoint management console, then log into the management console.

Step 7 Using the management console, authorize the AMP cloud to send Secure Endpoint data to Firewall Management Center.

Step 8 If you want to restrict the data that the Firewall Management Center receives, select specific groups within your organization for which you want to receive information.

By default, the AMP cloud sends data for all groups. To manage groups, choose **Management** > **Groups** on the Secure Endpoint management console. For detailed information, see the management console online help.

Step 9 Click **Allow** to enable the connection and start the transfer of data.

Clicking **Deny** returns you to the Firewall Management Center, where the connection is marked as denied. If you navigate away from the Applications page on the Secure Endpoint management console, and neither deny nor allow the connection, the connection is marked as pending on the Firewall Management Center's web interface. The health monitor does **not** alert you of a failed connection in either of these situations. If you want to connect to the AMP cloud later, delete the failed or pending connection, then recreate it.

Incomplete registration of the Secure Endpoint connection does not disable the malware defense connection.

Step 10 To verify that the connection is correctly configured:

a) On the **Integration** > **AMP** > **AMP Management** page, click the Cloud Name that includes Secure Endpoint in the **Cisco AMP Solution Type** column.
b) In the Secure Endpoint console window that displays, choose **Accounts > Applications**.
c) Verify that your Firewall Management Center is on the list.
d) In the Secure Endpoint console window, choose **Manage > Computers**.
e) Verify that your Firewall Management Center is on the list.

What to do next

- In the Secure Endpoint console window, configure settings as needed. For example, define group membership for your management center and assign policies. For information, see the Secure Endpoint online help or other documentation.
- In high availability deployments, configure AMP cloud connections on both management centers. These configurations are not synchronized.
- The default health policy warns you if the Firewall Management Center cannot connect to the Secure Endpoint portal after an initial successful connection, or if the connection is deregistered using the AMP portal.

  Verify that the **Secure Endpoint Status** monitor is enabled under **System > Health > Policy**.

History for Network Malware Protection and File Policies

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Communication with AMP cloud	7.0	Any	Legacy port 32137 is no longer supported for communicating with the AMP public or private cloud. New/Modified screens: On the **System** > **Integration** > **Cloud Services** page, the **Use Legacy Port 32137 for AMP for Networks** option is no longer available.
Chapter restructure	Changes were made in 6.4 timeframe, but will appear in any version that is republished	Any	Restructured this chapter's content to reduce confusion. Some content was moved to or from the chapter for *File/Malware Events and Network File Trajectory* in the Cisco Secure Firewall Management Center Administration Guide .
Moved URL Filtering information to the new URL Filtering chapter	6.3	Any	Moved information about configuring cloud communications for URL Filtering to the new URL Filtering chapter. Made related changes to the structure of the Cisco CSI topics in the chapter.

PART XI

Encrypted Traffic Handling

- Traffic Decryption Overview, on page 2253
- Decryption Policies, on page 2279
- Decryption Rules, on page 2303
- Decryption Rules and Policy Example, on page 2351

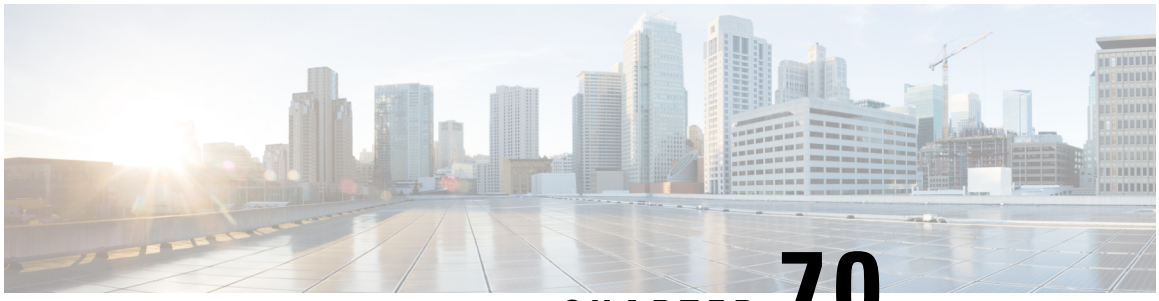

CHAPTER 70

Traffic Decryption Overview

The following topics provide an overview of Transport Layer Security/Secure Sockets Layer (TLS/SSL) inspection, discuss the prerequisites for TLS/SSL inspection configuration, and detail deployment scenarios.

Note Because TLS and SSL are often used interchangeably, we use the expression *TLS/SSL* to indicate that either protocol is being discussed. The SSL protocol has been deprecated by the IETF in favor of the more secure TLS protocol, so you can usually interpret *TLS/SSL* as referring to TLS only.

For more information about SSL and TLS protocols, see a resource such as SSL vs. TLS - What's the Difference?

- Traffic Decryption Explained, on page 2253
- TLS/SSL Handshake Processing, on page 2255
- Decryption Rule and Policy Basics, on page 2260
- TLS Crypto Acceleration, on page 2268
- DTLS Crypto Acceleration, on page 2271
- How to Configure Decryption Policies and Rules, on page 2272
- History for Decryption Policy, on page 2273

Traffic Decryption Explained

Most traffic on the internet is encrypted and in most cases, you might choose not to decrypt it. You can glean some information from encrypted traffic and you can also block it from your network if desired.

Your choices are:

- Decrypt the traffic and subject it to the full array of deep inspection:
  - Advanced Malware Protection
  - Security intelligence
  - Threat Intelligence Director
  - Application detectors
  - URL and category filtering

- Leave the traffic encrypted and set up your access control and decryption policy to look for and potentially block:
  - Old protocol versions (such as Secure Sockets Layer)
  - Unsecure cipher suites
  - Applications with high risk and low business relevance
  - Untrusted issuer Distinguished Names

Access control policy

An access control policy is the main configuration that invokes subpolicies and other configurations, including a decryption policy. If you associate a decryption policy with access control, the system uses that decryption policy to handle encrypted sessions before it evaluates the sessions with access control rules. If you do not configure TLS/SSL inspection, or your devices do not support it, access control rules handle all encrypted traffic.

Access control rules also handle encrypted traffic when your TLS/SSL inspection configuration allows the traffic to pass. However, some access control rule conditions require unencrypted traffic, so encrypted traffic might match fewer rules. Also, by default, the system disables intrusion and file inspection of encrypted payloads. This helps reduce false positives and improves performance when an encrypted connection matches an access control rule that has intrusion and file inspection configured.

Among the protocols we decrypt

Among the protocols we decrypt are the following:

- TLS versions up to 1.3 (if Snort 3 is enabled)
- DNS over HTTPS
- DNS over TLS (RFC 7858)
- FTPS, SMTPS, IMAPS, POP3S

Notes

Decryption rules require processing overhead that can impact performance. Before you decide to decrypt traffic, see When to Decrypt Traffic, When Not to Decrypt, on page 2261.

As long as your managed devices have Snort 3 enabled, the system supports decrypting TLS 1.3 traffic. You can enable TLS 1.3 decryption in an decryption policy's advanced options; for more information, see Decryption Policy Advanced Options, on page 2298.

The Firepower System does not support mutual authentication; that is, you cannot upload a client certificate to the Secure Firewall Management Center and use it for **Decrypt - Resign**, or **Decrypt - Known Key** decryption rule actions.

If you set the value of TCP maximum segment size (MSS) using FlexConfig, the observed MSS could be less than your setting. For more information, see About the TCP MSS, on page 878.

Related Topics

TLS/SSL Handshake Processing, on page 2255
Decryption Rule and Policy Basics, on page 2260

TLS/SSL Handshake Processing

In this documentation, the term *TLS/SSL handshake* represents the two-way handshake that initiates encrypted sessions in both the SSL protocol and its successor protocol, TLS.

In an inline deployment, the system processes the TLS/SSL handshake, potentially modifying the ClientHello message and acting as a TCP proxy server for the session.

The following figure shows an inline deployment.

After the client establishes a TCP connection with the server (after it successfully completes the TCP three-way handshake), the managed device monitors the TCP session for any attempt to initiate an encrypted session. The TLS/SSL handshake establishes an encrypted session using the exchange of specialized packets between client and server. In the SSL and TLS protocols, these specialized packets are called *handshake messages*. The handshake messages communicate which encryption attributes both the client and server support:

- ClientHello—The client specifies multiple supported values for each encryption attribute.
- ServerHello—The server specifies a single supported value for each encryption attribute, and the ServerHello response determines which encryption method the system uses during the secure session.

After a TLS/SSL handshake completes, the managed device caches encrypted session data, which allows session resumption without requiring the full handshake. The managed device also caches server certificate data, which allows faster handshake processing in subsequent sessions that use the same certificate.

ClientHello Message Handling

The client sends the ClientHello message to the server that acts as the packet destination if a secure connection can be established. The client sends the message to initiate the TLS/SSL handshake or in response to a ServerHello message from the destination server.

Overview

The following figure shows an example. Also see RFC 8446, sec. 4 . You can also consult a resource such as What Happens in a TLS Handshake? at cloudflare.com.

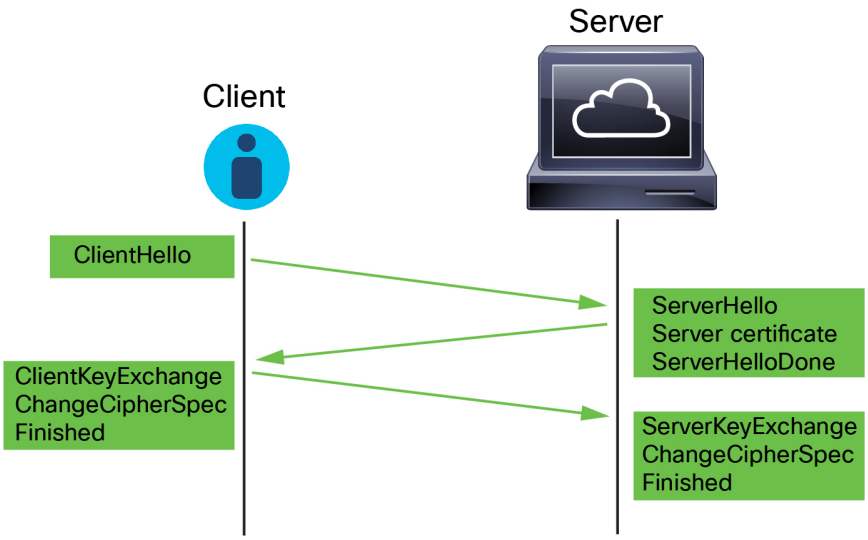

The process can be summarized as follows:

1. ClientHello initiates the process.

   The ClientHello message contains the Server Name Indication (SNI), which has the server's fully qualified domain name.

2. After a managed device processes a ClientHello message and transmits it to the destination server, the server determines whether it supports the decryption attributes the client specified in the message. If it does not support those attributes, the server sends a handshake failure alert to the client. If it supports those attributes, the server sends the ServerHello message. If the agreed-upon key exchange method uses certificates for authentication, the server certificate message immediately follows the ServerHello message.

   The server certificate contains the Subject Alternative Name (SAN), which can have fully qualified domain names and IP addresses. For more information about the SAN, see Distinguished Name, on page 1382.

3. When the managed device receives these messages, it attempts to match them with decryption rules configured on the system. These messages contain information that was absent from either the ClientHello message or the session data cache. Specifically, the system can potentially match these messages on decryption rules' Distinguished Names, Certificate Status, Cipher Suites, and Versions conditions.

The entire process is encrypted.

Data exchange

If you configure TLS/SSL decryption, when a managed device receives a ClientHello message, the system attempts to match the message to decryption rules that have the **Decrypt - Resign**, or **Decrypt - Known Key** action. The match relies on data from the ClientHello message and from cached server certificate data. Possible data includes:

Table 123: Data Availability for Decryption Rule Conditions

Decryption Rule Condition	Data Present In
Zones	ClientHello
Networks	ClientHello

Decryption Rule Condition	Data Present In
VLAN Tags	ClientHello
Ports	ClientHello
Users	ClientHello
Applications	ClientHello (Server Name Indicator extension)
Categories	ClientHello (Server Name Indicator extension)
Certificate	Server certificate (potentially cached)
Distinguished Names	Server certificate (potentially cached)
Certificate Status	Server certificate (potentially cached)
Cipher Suites	ServerHello
Versions	ServerHello

Important Use the **Cipher Suite** and **Version** rule conditions *only* in rules with either the **Block** or **Block with reset** rule actions. Do not use **Cipher Suite** and **Version** with **Decrypt - Resign** or **Decrypt - Known Key** rule actions. These conditions in rules with other rule actions can interfere with the system's ClientHello processing, resulting in unpredictable performance.

ClientHello modifications

If the ClientHello message matches a **Decrypt - Resign**, or **Decrypt - Known Key** rule, the system modifies the ClientHello message as follows:

- (TLS 1.2 only; TLS 1.3 does not support compression.) Compression methods—Strips the `compression_methods` element, which specifies the compression methods the client supports. The system cannot decrypt compressed sessions.
- Cipher suites—Strips cipher suites from the `cipher_suites` element if the system does not support them. If the system does not support any of the specified cipher suites, the system transmits the original, unmodified element. This modification reduces the Unknown Cipher Suite and Unsupported Cipher Suite types of undecryptable traffic.
- Session identifiers—Strips any value from the `Session Identifier` element and the SessionTicket extension (RFC 5077, sec 3.2) that does not match cached session data. If a ClientHello value matches cached data, an interrupted session can resume without the client and server performing the full TLS/SSL handshake. This modification increases the chances of session resumption and reduces the Session Not Cached type of undecryptable traffic.
- Elliptic curves—Strips elliptic curves from the Supported Elliptic Curves extension if the system does not support them. If the system does not support any of the specified elliptic curves, the managed device removes the extension and strips any related cipher suites from the `cipher_suites` element.
- ALPN extensions—Strips any value from the Application-Layer Protocol Negotiation (ALPN) extension that is unsupported in the system (for example, the HTTP/2 protocol).

- Other Extensions—Strips the Next Protocol Negotiation (NPN) and TLS Channel IDs extensions.

  Decryption rules with a **Decrypt - Resign**, or **Decrypt - Known Key** action now natively support the Extended Master Secret (EMS) extension during ClientHello negotiation, enabling more secure communications. The EMS extension is defined by RFC 7627.

After the system modifies the ClientHello message, it determines whether the message passes access control evaluation (which can include deep inspection). If the message passes, the system transmits it to the destination server.

If the ClientHello message does *not* match a **Decrypt - Resign**, **Decrypt - Replace Cert** or **Decrypt - Known Key** rule, the system does not modify the message. It then determines whether the message passes access control evaluation (which can include deep inspection). If the message passes inspection, the system transmits it to the destination server.

ClientHello is *not* modified if traffic matches a **Monitor** rule condition.

Man-in-the-middle

Direct communication between the client and server is no longer possible during the TLS/SSL handshake, because after message modification the Message Authentication Codes (MACs) computed by the client and server no longer match. For all subsequent handshake messages (and for the encrypted session once established), the managed device acts as a man-in-the-middle. It creates two TLS/SSL sessions, one between client and managed device, one between managed device and server. As a result, each session contains different cryptographic session details.

Note The cipher suites that the system can decrypt are frequently updated and do not correspond directly to the cipher suites you can use in decryption rule conditions. For the current list of decryptable cipher suites, contact Cisco TAC.

Related Topics

Default Handling Options for Undecryptable Traffic, on page 2296
ServerHello and Server Certificate Message Handling, on page 2258

ServerHello and Server Certificate Message Handling

Overview

The following figure shows an example. Also see RFC 8446, sec. 4 . You can also consult a resource such as What Happens in a TLS Handshake? at cloudflare.com.

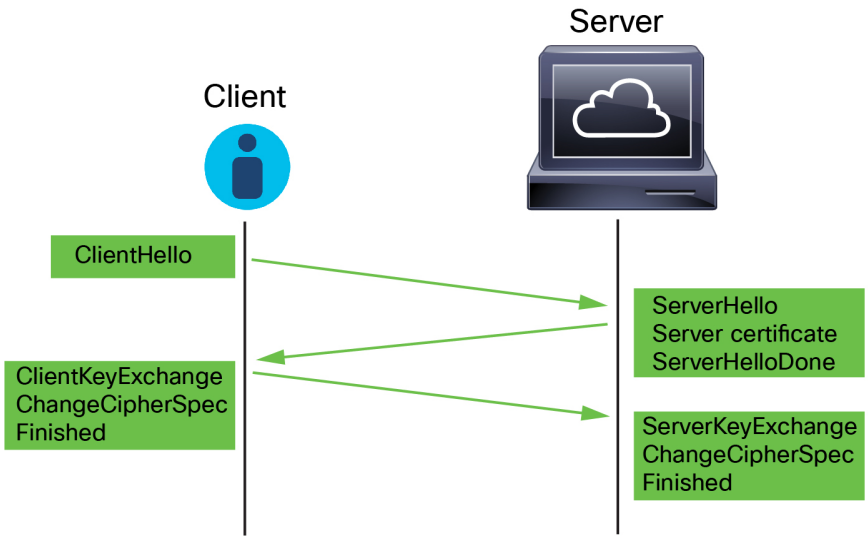

The process can be summarized as follows:

1. ClientHello initiates the process.

   The ClientHello message contains the Server Name Indication (SNI), which has the server's fully qualified domain name.

2. After a managed device processes a ClientHello message and transmits it to the destination server, the server determines whether it supports the decryption attributes the client specified in the message. If it does not support those attributes, the server sends a handshake failure alert to the client. If it supports those attributes, the server sends the ServerHello message. If the agreed-upon key exchange method uses certificates for authentication, the server certificate message immediately follows the ServerHello message.

   The server certificate contains the Subject Alternative Name (SAN), which can have fully qualified domain names and IP addresses. For more information about the SAN, see Distinguished Name, on page 1382.

3. When the managed device receives these messages, it attempts to match them with decryption rules configured on the system. These messages contain information that was absent from either the ClientHello message or the session data cache. Specifically, the system can potentially match these messages on decryption rules' Distinguished Names, Certificate Status, Cipher Suites, and Versions conditions.

The entire process is encrypted.

Decryption Rule actions

If the messages do not match any decryption rules, the managed device performs the Decryption Policy Default Actions, on page 2295.

If the messages match a rule that belongs to a decryption policy associated with an access control policy, the managed device continues as appropriate:

Action: Monitor
The TLS/SSL handshake continues to completion. The managed device tracks and logs traffic but does not decrypt it.

Action: Block or Block with Reset
The managed device blocks the TLS/SSL session and, if configured, resets the TCP connection.

Action: Do Not Decrypt

The TLS/SSL handshake continues to completion. The managed device does not decrypt the application data exchanged during the TLS/SSL session.

Action: Decrypt - Known Key

The managed device attempts to match the server certificate data to an Internal Certificate object previously imported into the Secure Firewall Management Center . Because you cannot generate an Internal Certificate object, and you must possess its private key, we assume you own the server on which you're using known key decryption.

If the certificate matches a known certificate, the TLS/SSL handshake continues to completion. The managed device uses the uploaded private key to decrypt and re-encrypt the application data exchanged during the TLS/SSL session.

If the server changes its certificate between the initial connection with the client and subsequent connections, you must import the new server certificate in the Secure Firewall Management Center for future connections to be decrypted.

Action: Decrypt - Resign

The managed device processes the server certificate message and re-signs the server certificate with the previously imported or generated certificate authority (CA). The TLS/SSL handshake continues to completion. The managed device then uses the uploaded private key to decrypt and re-encrypt the application data exchanged during the TLS/SSL session.

Note The Firepower System does not support mutual authentication; that is, you cannot upload a client certificate to the Secure Firewall Management Center and use it for **Decrypt - Resign**, or **Decrypt - Known Key** decryption rule actions.

Related Topics

ClientHello Message Handling, on page 2255

Decryption Rule and Policy Basics

This section discusses information you should keep in mind when creating your decryption policies and rules.

Note Because TLS and SSL are often used interchangeably, we use the expression *TLS/SSL* to indicate that either protocol is being discussed. The SSL protocol has been deprecated by the IETF in favor of the more secure TLS protocol, so you can usually interpret *TLS/SSL* as referring to TLS only.

For more information about SSL and TLS protocols, see a resource such as SSL vs. TLS - What's the Difference?

Related Topics

The Case for Decryption, on page 2261
When to Decrypt Traffic, When Not to Decrypt, on page 2261
Other Decryption Rule Actions, on page 2264

Decryption Rule Components, on page 2264
Decryption Rule Order Evaluation, on page 2265
TLS 1.3 Decryption Best Practices, on page 2300

The Case for Decryption

Traffic that is encrypted when it passes through the system can be allowed or blocked, but it *cannot* be subjected to deep inspection or the full range of policy enforcement (such as intrusion prevention).

All encrypted connections:

- Are sent through the decryption policy to determine if they should be decrypted or blocked.

  You can also configure decryption rules to block encrypted traffic of types you know you do not want on your network, such as traffic that uses the nonsecure SSL protocol or traffic with an expired or invalid certificate.

- If unblocked, whether or not decrypted, traffic goes through the access control policy for a final allow or block decision.

Only *decrypted* traffic takes advantage of the system's threat defense and policy enforcement features, including:

- Advanced Malware Protection
- Security intelligence
- Threat Intelligence Director
- Application detectors
- URL and category filtering

Keep in mind that decrypting and then re-encrypting traffic adds a processing load on the device, which can reduce overall system performance.

In summary:

- Encrypted traffic can be allowed or blocked by policy; encrypted traffic *cannot* be inspected
- Decrypted traffic is subject to threat defense and policy enforcement; decrypted traffic can be allowed or blocked by policy

Related Topics

Deep Inspection Using File and Intrusion Policies, on page 1728

When to Decrypt Traffic, When Not to Decrypt

This section provides guidelines on when you should decrypt traffic and when you should allow it to pass through the firewall encrypted.

When not to decrypt traffic

You should not decrypt traffic if doing so is forbidden by:

- Law; for example, some jurisdictions forbid decrypting financial information

- Company policy; for example, your company might forbid decrypting privileged communications
- Privacy regulations
- Traffic that uses certificate pinning (also referred to as *TLS/SSL pinning*) must remain encrypted to prevent breaking the connection

(Snort 3.) Decryption policy is *not* bypassed for any connections that match access control rules with actions of Trust, Block, or Block with Reset, unless the traffic is prefiltered. The encrypted traffic is first evaluated by decryption policy and then proceeds to the access control policy, where a final allow or block decision is made.

Encrypted traffic can be allowed or blocked on any decryption rule condition, including, but not limited to:

- Certificate status (for example, expired or invalid certificate)
- Protocol (for example, the nonsecure SSL protocol)
- Network (security zone, IP address, VLAN tag, and so on)
- URL category and reputation
- Port
- User group

Decryption rules provide a **Do Not Decrypt** action for traffic you do not want to decrypt; for more information, see Decryption Rule Do Not Decrypt Action, on page 2335.

Note The related information links at the end of this topic explain how some aspects of rule evaluation work. Conditions such as URL and application filtering have limitations with respect to encrypted traffic. Make sure you understand those limitations.

For more information about using URL filtering in **Do Not Decrypt** rules, see Decryption Rule Do Not Decrypt Action, on page 2335.

When to decrypt traffic

All encrypted traffic must be decrypted to take advantage of the system's threat protection and policy enforcement features. If your managed device has sufficient memory and processing power to decrypt traffic, you should decrypt traffic that is not prohibited by law or regulation.

If you must decide what traffic to decrypt, base your decision on the risk of allowing the traffic on your network. The system provides a flexible framework for classifying traffic using rule conditions, which include URL reputation, cipher suite, protocol, and many other factors.

Related Topics

Decrypt and Resign (Outgoing Traffic), on page 2263
Known Key Decryption (Incoming Traffic), on page 2263
Decryption Rule Guidelines and Limitations, on page 2304
Decryption Policy Rule Order
URL Conditions (URL Filtering)
Application Rule Order, on page 1741

TLS 1.3 Decryption Best Practices, on page 2300

Decrypt and Resign (Outgoing Traffic)

This information applies only to rule-based decryption policies and rules.

The **Decrypt - Resign** decryption rule action enables the system to act as a man in the middle, intercepting, decrypting, and (if the traffic is allowed to pass) inspecting and re-encrypting it. The **Decrypt - Resign** rule action is used with outgoing traffic; that is, the destination server is outside your protected network.

The Firewall Threat Defense device negotiates with the client using an internal Certificate Authority (CA) object specified in the rule and builds a TLS/SSL tunnel between the client and the Firewall Threat Defense device. At the same time, the device connects to the destination web site and creates a TLS/SSL tunnel between the server and the Firewall Threat Defense device.

Thus, the client sees the CA certificate configured for the decryption rule instead of the certificate from the destination server. The client must trust the firewall's certificate to complete the connection. The Firewall Threat Defense device then performs decryption/re-encryption in both directions for traffic between the client and the destination server.

Prerequisite

To use the **Decrypt - Resign** rule action, you must create an internal CA object using a CA file and paired private key file. You can generate a CA and private key in the system if you don't already have them.

We simplify the process of configuring the **Decrypt - Resign** rule by automatically adding **Do Not Decrypt** rules to the decryption policy based on your choices for types of traffic to exempt from encryption, such as undecryptable applications. (Typically, these applications use certificate pinning so that decrypting the connection breaks the connection.) For more information, see Create a Create a Decryption Policy with Outbound Connection Protection, on page 2281.

Note The Firepower System does not support mutual authentication; that is, you cannot upload a client certificate to the Secure Firewall Management Center and use it for **Decrypt - Resign**, or **Decrypt - Known Key** decryption rule actions.

Related Topics

Decryption Rule Decrypt Actions, on page 2337
External Certificate Objects, on page 1411

Known Key Decryption (Incoming Traffic)

This information applies only to rule-based decryption policies and rules.

The **Decrypt - Known Key** decryption rule action uses a server's private key to decrypt traffic. The **Decrypt - Known Key** rule action is used with incoming traffic; that is, the destination server is inside your protected network.

The main purpose of decrypting with a known key is to protect your servers from external attacks.

Prerequisite

To use the **Decrypt - Known Key** rule action, you must create an internal certificate object using the server's certificate file and paired private key file.

Note The Firepower System does not support mutual authentication; that is, you cannot upload a client certificate to the Secure Firewall Management Center and use it for **Decrypt - Resign**, or **Decrypt - Known Key** decryption rule actions.

Related Topics

Known Key Decryption (Incoming Traffic), on page 2263
Decryption Rule Decrypt Actions, on page 2337
Internal Certificate Objects, on page 1412

Other Decryption Rule Actions

This information applies only to rule-based decryption policies and rules.

The following sections discuss other decryption rule actions.

Related Topics

Decryption Rule Blocking Actions, on page 2336
Decryption Rule Monitor Action, on page 2334

Decryption Rule Components

This information applies only to rule-based decryption policies and rules.

Each decryption rule has the following components.

State

A rule can have either of two states: enabled or disabled. By default, rules are enabled. If you disable a rule, the system does not use it to evaluate network traffic, and stops generating warnings and errors for that rule.

Position

Rules in a decryption policy are numbered, starting at 1. The system matches traffic to rules in top-down order by ascending rule number.

Conditions

Conditions specify the specific traffic the rule handles. Conditions can match traffic by security zone, network or geographical location, VLAN, port, application, requested URL category and reputation, user, certificate, certificate subject or issuer, certificate status, cipher suite, or encryption protocol version. The use of conditions can depend on target device licenses.

Action

A rule's action determines how the system handles matching traffic. You can monitor, allow, block, or decrypt encrypted matching traffic. Decrypted traffic is subject to further inspection. Note that the system does *not* inspect blocked encrypted traffic.

Logging

A rule's logging settings govern the records the system keeps of the traffic it handles. You can keep a record of traffic that matches a rule. You can log a connection when the system blocks an encrypted session or allows it to pass without decryption, according to the settings in an decryption policy. You can also force the system to log connections that it decrypts for further evaluation by access control rules, regardless of how the system later handles or inspects the traffic. You can log connections to the Secure Firewall Management Center database, as well as to the system log (syslog) or to an SNMP trap server.

For more information about logging, see Best Practices for Connection Logging in the Cisco Secure Firewall Management Center Administration Guide.

Tip Properly creating and ordering decryption rules is a complex task. If you do not plan your policy carefully, rules can preempt other rules, require additional licenses, or contain invalid configurations. To help ensure that the system handles traffic as you expect, the decryption policy interface has a robust warning and error feedback system for rules.

Categories

For information about using decryption rule categories (such as Applications, Category, and Cert Status), see Decryption Rule Conditions, on page 2316.

Decryption Rule Order Evaluation

This information applies only to rule-based decryption policies and rules.

When you create a decryption rule in a decryption policy, you specify its position using the **Insert** list in the rule editor. Decryption rules in an a decryption policy are numbered, starting at 1. The system matches traffic to decryption rules in top-down order by ascending rule number.

In most cases, the system handles network traffic according to the *first* decryption rule where *all* the rule's conditions match the traffic. Except in the case of Monitor rules (which log traffic but do not affect traffic flow), the system does *not* continue to evaluate traffic against additional, lower-priority rules after that traffic matches a rule. Conditions can be simple or complex; you can control traffic by security zone, network or geographical location, VLAN, port, application, requested URL category and reputation, user, certificate, certificate distinguished name, certificate status, cipher suite, or encryption protocol version.

Each rule also has an *action*, which determines whether you monitor, block, or inspect matching encrypted or decrypted traffic with access control. Note that the system does *not* further inspect encrypted traffic it blocks. It does subject encrypted and undecryptable traffic to access control. However, access control rule conditions require unencrypted traffic, so encrypted traffic matches fewer rules.

Rules that use *specific* conditions (such as network and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your rules. For more information about the OSI model, see this Wikipedia article.

Tip Proper decryption rule order reduces the resources required to process network traffic, and prevents rule preemption. Although the rules you create are unique to every organization and deployment, there are a few general guidelines to follow when ordering rules that can optimize performance while still addressing your needs.

In addition to ordering rules by number, you can group rules by category. By default the system provides three categories: Administrator, Standard, and Root. You can add custom categories, but you cannot delete the system-provided categories or change their order.

Related Topics

Best Practices for Access Control Rules, on page 1738
Default Handling Options for Undecryptable Traffic, on page 2296
Decryption Policy Rule Order

Multi-Rule Example

The following scenario summarizes the ways that decryption rules handle traffic in an inline deployment.

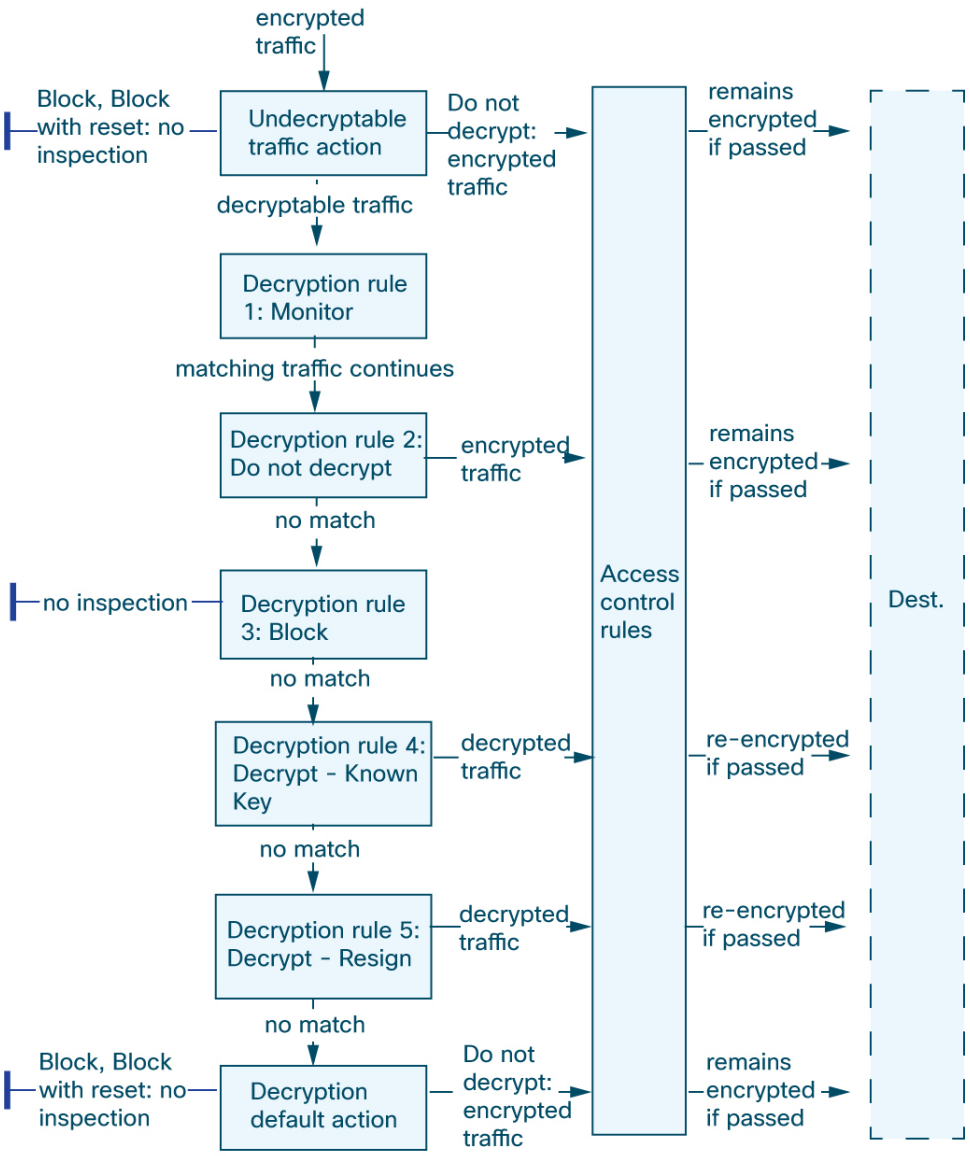

In this scenario, traffic is evaluated as follows:

- **Undecryptable Traffic Action** evaluates encrypted traffic first. For traffic the system cannot decrypt, the system either blocks it without further inspection or passes it for access control inspection. Encrypted traffic that does not match continues to the next rule.
- **Decryption Rule 1: Monitor** evaluates encrypted traffic next. Monitor rules track and log encrypted traffic but do not affect traffic flow. The system continues to match traffic against additional rules to determine whether to permit or deny it.
- **Decryption Rule 2: Do Not Decrypt** evaluates encrypted traffic third. Matching traffic is not decrypted; the system inspects this traffic with access control, but not file or intrusion inspection. Traffic that does not match continues to the next rule.

- **Decryption Rule 3: Block** evaluates encrypted traffic fourth. Matching traffic is blocked without further inspection. Traffic that does not match continues to the next rule.
- **Decryption Rule 4: Decrypt - Known Key** evaluates encrypted traffic fifth. Matching traffic incoming to your network is decrypted using a private key you upload. The decrypted traffic is then evaluated against access control rules. Access control rules handle decrypted and unencrypted traffic identically. The system can block traffic as a result of this additional inspection. All remaining traffic is reencrypted before being allowed to the destination. Traffic that does not match the decryption rule continues to the next rule.
- **Decryption Rule 5: Decrypt - Resign** is the final rule. If traffic matches this rule, the system re-signs the server certificate with an uploaded CA certificate, then acts as a man-in-the-middle to decrypt traffic. The decrypted traffic is then evaluated against access control rules. Access control rules treat decrypted and unencrypted traffic identically. The system can block traffic as a result of this additional inspection. All remaining traffic is reencrypted before being allowed to the destination. Traffic that does not match the decryption rule rule continues to the next rule.
- **Decryption policy Default Action** handles all traffic that does not match any of the decryption rules. The default action either blocks encrypted traffic without further inspection or does not decrypt it, passing it for access control inspection.

TLS Crypto Acceleration

TLS crypto acceleration accelerates the following:

- TLS/SSL encryption and decryption
- VPN, including TLS/SSL and IPsec

Supported Hardware

The following hardware models support TLS crypto acceleration:

- Secure Firewall 3100
- Secure Firewall 4200
- Firepower 4100/9300

  For information about TLS crypto acceleration support on Firepower 4100/9300 Firewall Threat Defense container instances, see the *FXOS Configuration Guide*.

TLS crypto acceleration is *not* supported on any virtual appliances or on any hardware except for the preceding.

Note For more information about TLS crypto acceleration and the 4100/9300, see the *FXOS Configuration Guide*.

Features Not Supported by TLS crypto acceleration

Features *not* supported by TLS crypto acceleration include the following:

- Managed devices where Firewall Threat Defense container instance is enabled.

- If the inspection engine is configured to preserve connections and the inspection engine fails unexpectedly, TLS/SSL traffic is dropped until the engine restarts.

  This behavior is controlled by the **configure snort preserve-connection {enable | disable}** command.

TLS Crypto Acceleration Guidelines and Limitations

Keep the following in mind if your managed device has TLS crypto acceleration enabled.

HTTP-only performance

Using TLS crypto acceleration on a managed device that is not decrypting traffic can affect performance.

Federal Information Processing Standards (FIPS)

If TLS crypto acceleration and Federal Information Processing Standards (FIPS) are both enabled, connections with the following options fail:

- RSA keys less than 2048 bytes in size
- Rivest cipher 4 (RC4)
- Single Data Encryption Standard (single DES)
- Merkle–Damgard 5 (MD5)
- SSL v3

FIPS is enabled when you configure the Firewall Management Center and managed devices to operate in a security certifications compliance mode. To allow connections when operating in those modes, you can configure web browsers to accept more secure options.

For more information:

- Ciphers supported by FIPS: About SSL Settings, on page 980.
- Security Certifications Compliance Modes.
- Common Criteria.

TLS heartbeat

Some applications use the *TLS heartbeat* extension to the Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) protocols defined by RFC6520. TLS heartbeat provides a way to confirm the connection is still alive—either the client or server sends a specified number of bytes of data and requests the other party echo the response. If this is successful, encrypted data is sent.

When a managed device with TLS crypto acceleration enabled encounters a packet that uses the TLS heartbeat extension, the managed device takes the action specified by the setting for **Decryption Errors** in the decryption policy's **Undecryptable Actions**:

- Block
- Block with reset

For more information, see Default Handling Options for Undecryptable Traffic, on page 2296.

To determine whether applications are using TLS heartbeat, see Troubleshoot TLS Heartbeat, on page 2342.

You can configure a **Max Heartbeat Length** in a Network Analysis Policy (NAP) to determine how to handle TLS heartbeats. For more information, see The SSL Preprocessor.

TLS/SSL oversubscription

TLS/SSL oversubscription is a state where a managed device is overloaded with TLS/SSL traffic. Any managed device can experience TLS/SSL oversubscription but only managed devices that support TLS crypto acceleration provide a configurable way to handle it.

When a managed device with TLS crypto acceleration enabled is oversubscribed, any packet received by the managed device is acted on according to the setting for **Handshake Errors** in the decryption policy's **Undecryptable Actions**:

- Inherit default action
- Do not decrypt
- Block
- Block with reset

If the setting for **Handshake Errors** in the decryption policy's **Undecryptable Actions** is **Do Not decrypt** and the associated access control policy is configured to inspect the traffic, inspection occurs; decryption does *not* occur.

If a significant amount of oversubscription is occurring, you have the following options:

- Upgrade your managed devices to increase TLS/SSL processing capacity.
- Change your decryption policies to add **Do Not Decrypt** rules for traffic that is not a high priority to decrypt.

View the Status of TLS Crypto Acceleration

This topic discusses how you can determine if TLS crypto acceleration is enabled.

Perform the following task in the Firewall Management Center.

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Devices** > **Device Management**.

Step 3 Click **Edit** (✎) to edit a managed device.

Step 4 Click **Device** page. TLS crypto acceleration status is displayed in the General section.

DTLS Crypto Acceleration

The Firewall Threat Defense, with the help of the FPGA and the Nitrox V crypto accelerator, supports DTLS cryptographic acceleration for the following models:

- Secure Firewall 3100
- Secure Firewall 4200

This feature improves the throughput of the DTLS-encrypted and DTLS-decrypted traffic. Both IPv4 and IPv6 traffic are supported.

The Firewall Threat Defense also performs optimization of the egress-encrypted packets to improve latency. The data path is optimized to enhance performance for single tunnel flows.

Both features are enabled by default and work only for DTLS 1.2.

These features are not supported with:

- Flows that use DTLS 1.0 or packet compression
- Rekeyed DTLS keys
- Clustering or multi-instance mode

You can disable DTLS crypto acceleration using FlexConfig and the **no flow-offload-dtls** and the **no flow-offload-dtls egress-optimization** commands..

Prerequisites for DTLS Crypto Acceleration

Supported License

- Management center Essentials (formerly Base) license must allow export-controlled functionality.

  Choose **System > Licenses > Smart Licenses** to verify this functionality in the management center.

Supported Devices

- Secure Firewall 3100
- Secure Firewall 4200

Monitoring DTLS Crypto Acceleration

Use the following CLI commands on the threat defense device to verify and monitor DTLS crypto acceleration and optimization of egress-encrypted packets.

- To verify the status of DTLS crypto acceleration and optimization of egress-encrypted packets, use the following command:

```
> show flow-offload-dtls info
  DTLS offload : Enabled
```

```
  Egress Optimization: Enabled
```

- To view the DTLS crypto acceleration statistics, use the following command:

```
> show flow-offload-dtls statistics
  Packet stats of Pipe 0
  ----------------------
  Rx Packet count : 975638666
  Tx Packet count : 975638666
  Error Packet count : 0
  Drop Packet count : 0

  CAM stats of Pipe 0
  ----------------------
  Option ID Table CAM Hit Count : 1145314723
  Option ID Table CAM Miss Count : 0
  Tunnel Table CAM Hit Count : 0
  Tunnel Table CAM Miss Count : 0
  6-Tuple CAM Hit Count : 975638666
  6-Tuple CAM Miss Count : 169676057
   NOTE: The counters displayed are cumulative counters
   for all offload applications and indicates the total packets
  offloaded
```

- To view the device's Nitrox V crypto accelerator statistics, use the following command:

```
> show crypto accelerator statistics

Crypto Accelerator Status
-------------------------
<snip>
[Offloaded SSL Input statistics, Pipe 0]
      Input packets: 290593023
      Input bytes: 147049729714
      Decrypted packets: 290593023
      Decrypted bytes: 147049729714
   [Offloaded SSL Output statistics, Pipe 0]
      Output packets: 254271808
      Output bytes: 136352952720
      Encrypted packets: 254271808
      Encrypted bytes: 136352952720
.
.
.
```

How to Configure Decryption Policies and Rules

This topic provides a high-level overview of tasks you must complete to configure decryption policies and decryption rules in those policies to block, monitor, or allow TLS/SSL traffic on your network.

You must be an Admin, Access Admin, or Network Admin to perform this task.

Procedure

Step 1 (Optional.) For **Decrypt - Known Key** decryption rules (to decrypt inbound traffic to an internal server), create an internal certificate object.

The internal certificate object uses your server's certificate and private key. See Internal Certificate Objects, on page 1412.

Step 2 (Optional.) For **Decrypt - Resign** decryption rules (to decrypt outbound traffic to a server outside of your network), create an internal certificate authority (CA) object.
The internal CA object uses a CA and private key. See Internal Certificate Authority Objects, on page 1405.

Step 3 For outbound protection policies (**Decrypt - Resign**) or inbound protection policies (**Decrypt - Known Key**), run the decryption policy wizard.

- Create a Create a Decryption Policy with Outbound Connection Protection, on page 2281
- Create a Create a Decryption Policy with Inbound Connection Protection, on page 2285

Step 4 For any other rule action (**Do Not Decrypt**, **Block**, **Block with Reset**, or **Monitor**), create a decryption policy and rules manually.
See Create a Create a Decryption Policy with Other Rule Actions, on page 2294.

What to do next

See one of the following:

- Create a Create a Decryption Policy with Outbound Connection Protection, on page 2281
- Create a Create a Decryption Policy with Inbound Connection Protection, on page 2285
- Create a Create a Decryption Policy with Other Rule Actions, on page 2294

History for Decryption Policy

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Hardware DTLS 1.2 crypto acceleration for the Secure Firewall 3100/4200.	7.6.0	7.6.0	The Secure Firewall 3100/4200 now supports DTLS 1.2 cryptographic acceleration and egress optimization, which improves throughput of DTLS-encrypted and decrypted traffic. This is automatically enabled on new and upgraded devices. To disable, use FlexConfig. New/modified FlexConfig commands: **flow-offload-dtls**, **flow-offload-dtls egress-optimization**, **show flow-offload-dtls**

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
QUIC decryption.	7.6.0	7.6.0 with Snort 3	You can configure the decryption policy to apply to sessions running on the QUIC protocol. QUIC decryption is disabled by default. You can selectively enable QUIC decryption per decryption policy and write decryption rules to apply to QUIC traffic. By decrypting QUIC connections, the system can then inspect the connections for intrusion, malware, or other issues. You can also apply granular control and filtering of decrypted QUIC connections based on specific criteria in the access control policy. We modified the decryption policy Advanced Settings to include the option to enable QUIC decryption.
Easily bypass decryption for sensitive and undecryptable traffic.	7.6.0	7.6.0	It is now easier to bypass decryption for sensitive and undecryptable traffic, which protects users and improves performance. New decryption policies now include predefined rules that, if enabled, can automatically bypass decryption for sensitive URL categories (such as finance or medical), undecryptable distinguished names, and undecryptable applications. Distinguished names and applications are undecryptable typically because they use TLS/SSL certificate pinning, which is itself not decryptable. For outbound decryption, you enable/disable these rules as part of creating the policy. For inbound decryption, the rules are disabled by default. After the policy is created, you can edit, reorder, or delete the rules entirely. New/modified screens: **Policies** > **Access Control** > **Decryption** > **Create Decryption Policy**

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Decryption policy.	7.3.0	7.3.0	Feature renamed to *decryption policy* to better reflect what it does. We now enable you to configure a decryption policy with one or more **Decrypt - Resign** or **Decrypt - Known Key** rules at the same time. New/modified screens: • **Policies** > **Access Control heading** > **Decryption** (create new decryption policy) • The Create Decryption Policy dialog box now has two tab pages: **Outbound Connections** and **Inbound Connections**. Use the **Outbound Connections** tab page to configure one or more decryption rules with a **Decrypt - Resign** rule action. (You can either upload or generate certificate authorities at the same time). Each combination of a CA with networks and ports results in one decryption rule. Use the Inbound Connections tab page to configure one or more decryption rules with a Decrypt - Known Key rule action. (You can upload your server's certificate at the same time.) Each combination of a server certificate with networks and ports results in one decryption rule. • **Policies** > **Access Control heading** > **Decryption** (edit a decryption rule) **Advanced Settings** has new options discussed in TLS 1.3 Decryption Best Practices, on page 2300. • **Policies** > **Access Control Heading** > **Access Control** (edit an access control policy), click the word **Decryption** to associate a decryption policy with an access control policy.
TLS 1.3 decryption.	7.2.0	7.2.0	You can now enable TLS 1.3 decryption in an SSL policy's advanced actions. TLS 1.3 decryption requires the managed device run Snort 3. Other options are available as well; for more information, see TLS 1.3 Decryption Best Practices, on page 2300. New/modified screens: **SSL Policy** > **Advanced Settings**
SSL policy advanced settings.	7.2.0	7.1.0	SSL policy advanced settings New/modified screens: **SSL Policy** > **Advanced Settings**
Ability to specify handling of URLs having unknown reputation.	6.7.0	6.7.0	For details, see About URL Filtering with Category and Reputation, on page 1858.
ClientHello modification for **Decrypt - Known** key rules.	6.7.0	6.7.0	For details, see ClientHello Message Handling, on page 2255.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Ability to extract the certificate in TLS 1.3 traffic to enable traffic to match URL and application criteria in access control rules.	6.7.0	6.7.0	New/modified screens: **Policies** > **Access Control** > **(edit an access control policy)** > **Advanced** link. For details, see Decryption Policy Advanced Options, on page 2298.
Changes to category and reputation-based URL Filtering.	6.7.0	6.5.0	For details, see About URL Filtering with Category and Reputation, on page 1858.
TLS crypto acceleration cannot be disabled.	6.4.0	6.4.0	TLS crypto acceleration is enabled on all supported devices. On a managed device with native interfaces, TLS crypto acceleration cannot be disabled. Support for TLS crypto acceleration on Firewall Threat Defense container instances is limited as discussed in the next row of this table. Removed commands: • **system support ssl-hw-accel enable** • **system support ssl-hw-accel disable** • **system support ssl-hw-status**
Support for TLS crypto acceleration on one Firewall Threat Defense container instance on a Firepower 4100/9300 module/security engine.	6.4.0	6.4.0	You can now enable TLS crypto acceleration for one Firewall Threat Defense container instance on a module/security engine. TLS crypto acceleration is disabled for other container instances, but enabled for native instances. New/modified commands: • **config hwCrypto enable** • **show crypto accelerator status** replaces **system support ssl-hw-status**)
TLS/SSL hardware acceleration is now referred to as *TLS crypto acceleration*.	6.4.0	6.4.0	The name change reflects that TLS/SSL encryption and decryption acceleration is supported on more devices. Depending on the device, acceleration might be performed in software or in hardware. New/modified screens: **Devices** > **Device Management** > **Edit** > **Device** > **General** > **TLS Crypto Acceleration**
Extended Master Secret extension supported (see RFC 7627).	6.3.0.1	6.3.0.1	The TLS Extended Master Secret extension is supported for SSL policies; specifically, policies with a rule action of **Decrypt - Resign** or **Decrypt - Known Key**.
Extended Master Secret extension not supported.	6.3.0	6.3.0	The extension is stripped during ClientHello modification for **Decrypt - Resign** rules.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
TLS/SSL hardware acceleration enabled by default.	6.3.0	6.3.0	TLS/SSL hardware acceleration is enabled by default on all supported devices but can be disabled if desired.
Extended Master Secret extension supported (see RFC 7627).	6.2.3.9	6.2.3.9	The TLS Extended Master Secret extension is supported for SSL policies; specifically, policies with a rule action of **Decrypt - Resign** or **Decrypt - Known Key**.
Aggressive TLS 1.3 downgrade.	6.2.3.7	6.2.3.7	Using the **system support ssl-client-hello-enabled aggressive_tls13_downgrade {true|false}** CLI command, you can determine the behavior for downgrading TLS 1.3 traffic to TLS 1.2. For details, see the Cisco Secure Firewall Threat Defense Command Reference.
TLS/SSL hardware acceleration introduced.	6.2.3	6.2.3	Certain managed device models perform TLS/SSL encryption and decryption in hardware, improving performance. By default, the feature is enabled. Affected screen: To view the status of TLS/SSL hardware acceleration, **Devices** > **Device Management** > **Device**, General page.
Category and reputation conditions supported.	6.2.2	6.2.2	Access control rules or SSL rules with category/reputation conditions.
SafeSearch supported.	6.1.0	6.1.0	The system displays an HTTP response page for connections decrypted by the SSL policy, then blocked (or interactively blocked) either by access control rules or by the access control policy default action. In these cases, the system encrypts the response page and sends it at the end of the reencrypted SSL stream. SafeSearch filters objectionable content and stops people from searching adult sites.
TLS/SSL policy.	6.0.0	6.0.0	Feature introduced.

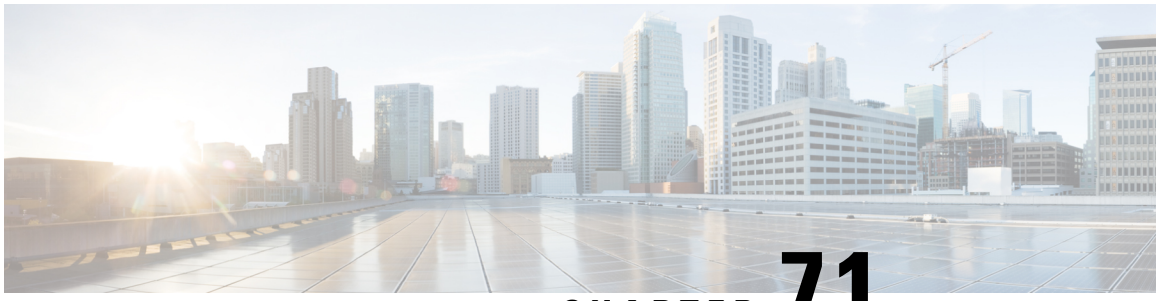

CHAPTER 71

Decryption Policies

The following topics provide an overview of decryption policy creation, configuration, management, and logging.

- About Decryption Policies, on page 2279
- Requirements and Prerequisites for Decryption Policies, on page 2281
- Create a Decryption Policy, on page 2281
- Decryption Policy Default Actions, on page 2295
- Default Handling Options for Undecryptable Traffic, on page 2296
- Decryption Policy Advanced Options, on page 2298

About Decryption Policies

A decryption policy determines how the system handles encrypted traffic on your network. You can configure one or more decryption policies, associate a decryption policy with an access control policy, then deploy the access control policy to a managed device. When the device detects a TCP handshake, the access control policy first handles and inspects the traffic. If it subsequently identifies a TLS/SSL-encrypted session over the TCP connection, the decryption policy takes over, handling and decrypting the encrypted traffic.

Create a decryption policy using the wizard

You can use a wizard to create the following types of decryption policies:

- Outbound protection (**Decrypt - Resign** rule action). If traffic matches this rule, the system re-signs the server certificate with the CA certificate, then acts as a man-in-the-middle.

  Three rules with a **Do Not Decrypt** action are added to the policy at the same time, saving you the trouble of doing it later. These rules correspond to any decryption exclusions you configure when you create the policy (for example, you can choose to bypass decryption for applications that are known to use certificate pinning.

  For more information, see Create a Create a Decryption Policy with Outbound Connection Protection, on page 2281.

- Inbound protection (**Decrypt - Known Key** rule action). You can associate one or more server certificates and paired private keys with the action. If traffic matches the rule, and the certificate used to encrypt the traffic matches the certificate associated with the action, the system uses the appropriate private key to obtain the session encryption and decryption keys.

Three rules with a **Do Not Decrypt** action are added to the policy at the same time but these rules are disabled by default. These rules correspond to any decryption exclusions you configure when you create the policy (for example, you can choose to bypass decryption for applications that are known to use certificate pinning.

For more information, see Create a Create a Decryption Policy with Inbound Connection Protection, on page 2285.

- Any other decryption rule action (such as blocking or monitoring).

  For more information, see Create a Create a Decryption Policy with Other Rule Actions, on page 2294.

The wizard automatically creates a separate rule for each certificate you specify. For example, an inbound protection rule might specify one certificate for traffic coming to the Finance department internal network and a different certificate for traffic coming to the Engineering network.

The wizard creates additional rules for outbound and inbound protection policies as follows:

- Outbound protection (**Decrypt - Resign** rule action): The wizard creates Do Not Decrypt rules for traffic that matches exceptions you specify in the wizard. For example, you can choose to not decrypt traffic from undecryptable applications, typically those using certificate pinning.

  The Do Not Decrypt rules are placed first in the decryption policy so that traffic is passed through the firewall with minimal processing.

- Inbound protection (**Decrypt - Known Key** rule action): The wizard does not allow you to choose any exceptions but it adds the Do Not Decrypt rules to the policy and disables them; that way, you can enable those exceptions later if you need to.

Do Not Decrypt policy example

Following is an example decryption policy with a **Do Not Decrypt** rule action:

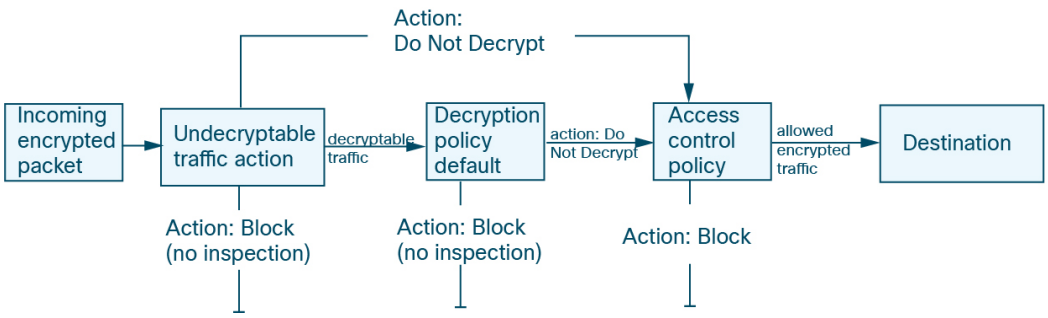

The simplest decryption policy, as shown in the following diagram, directs the device where it is deployed to handle encrypted traffic with a single default action. You can set the default action to block decryptable traffic without further inspection, or to inspect undecrypted decryptable traffic with access control. The system can then either allow or block the encrypted traffic. If the device detects undecryptable traffic, it either blocks the traffic without further inspection or does not decrypt it, inspecting it with access control.

To get started, see Create a Decryption Policy, on page 2281

Requirements and Prerequisites for Decryption Policies

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Create a Decryption Policy

You can create any of the following types of decryption policies:

- *Outbound protection* policy with rules that protect outbound connections; that is, the destination server is outside your protected network. This type of rule has a **Decrypt - Resign** rule action. We also create additional rules with a **Do Not Decrypt** action that excludes traffic you specify (such as traffic that uses certificate pinning).

  See Create a Create a Decryption Policy with Outbound Connection Protection, on page 2281

- *Inbound protection* policy with a rule that protects inbound connections; that is, the destination server is inside your protected network. This type of rule has a **Decrypt - Known Key** rule action. We also create additional rules with a **Do Not Decrypt** action that excludes traffic you specify (such as traffic that uses certificate pinning.) These rules are disabled initially but you can modify and enable them later if you wish.

  See Create a Create a Decryption Policy with Inbound Connection Protection, on page 2285

- Other actions (including **Do Not Decrypt**, **Block**, and **Block with Reset**).

  See Create a Create a Decryption Policy with Other Rule Actions, on page 2294

Create a Create a Decryption Policy with Outbound Connection Protection

This task discusses how to create a decryption policy with a rule that protects outbound connections; that is, the destination server is outside your protected network. This type of rule has a **Decrypt - Resign** rule action.

When you create a decryption policy, you can create multiple rules at the same time, including multiple **Decrypt - Known Key** rules, multiple **Decrypt - Replace Cert** rules, and multiple **Decrypt - Resign** rules.

If you enabled Change Management, you must create and assign a ticket before you can create a decryption policy. Before the decryption policy can be used, the ticket and all associated objects (like certificate authorities) must be approved. For more information, see Creating Change Management Tickets, on page 240 and Policies and Objects that Support Change Management, on page 235.

Before you begin

You can optionally must upload or generate an internal CA certificate for your managed device before you can create a decryption policy that protects outbound connections. You can do this in any of the following ways:

- Create an internal CA certificate object by going to **Objects** > **Object Management** > **PKI** > **Internal CAs** and referring to PKI, on page 1404.
- At the time you create this decryption policy.

Procedure

Step 1 Log in to Secure Firewall Management Center if you haven't already done so.

Step 2 Click **Policies** > **Access Control heading** > **Decryption**.

Step 3 Click **Create Decryption Policy**.

Step 4 Give the policy a unique **Name** and, optionally, a **Description**.

The following characters are not supported in decryption policy names:

- Leading period
- `#,;,{,},=,$,<,>`

Step 5 Click the **Outbound Connections** tab.

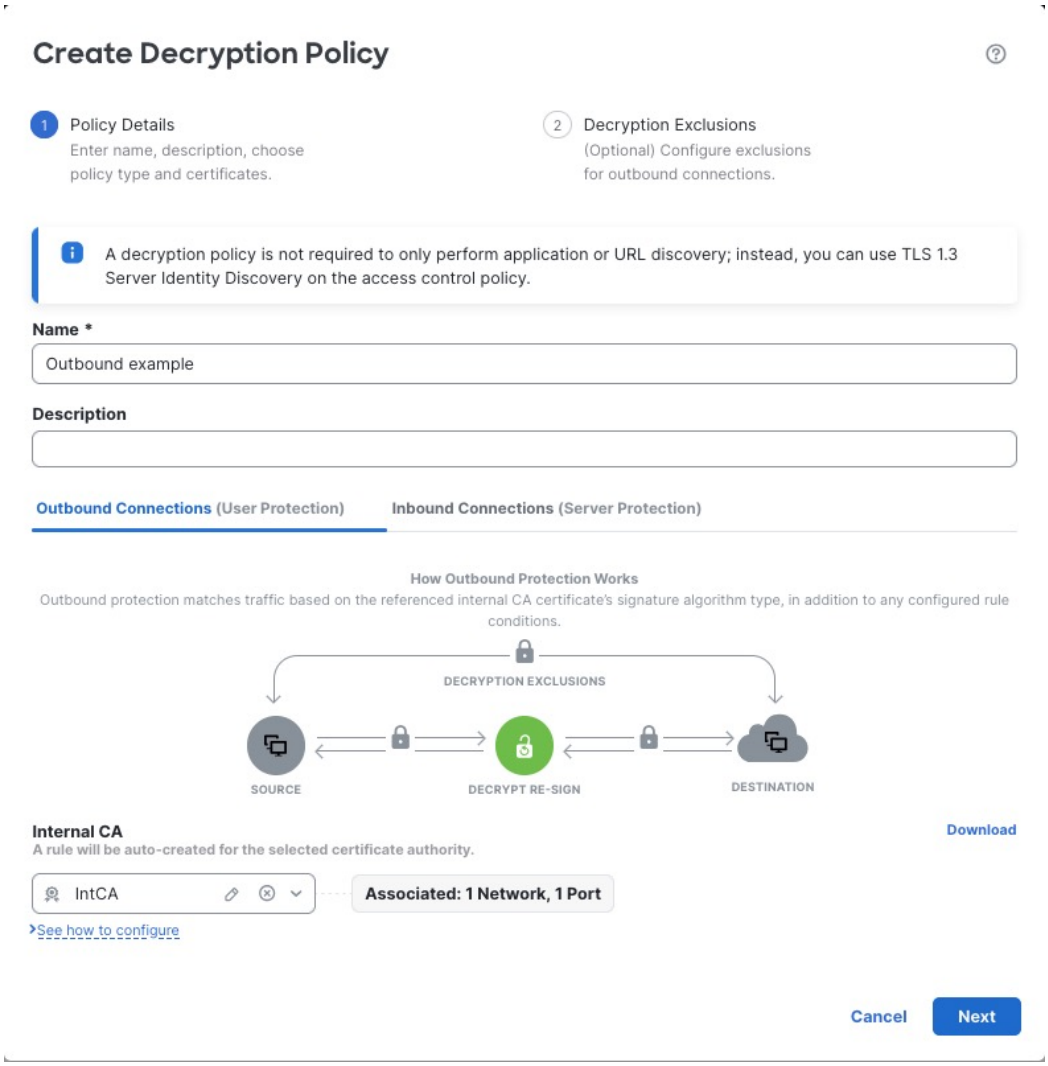

Step 6 Click the **Outbound Connections** tab.

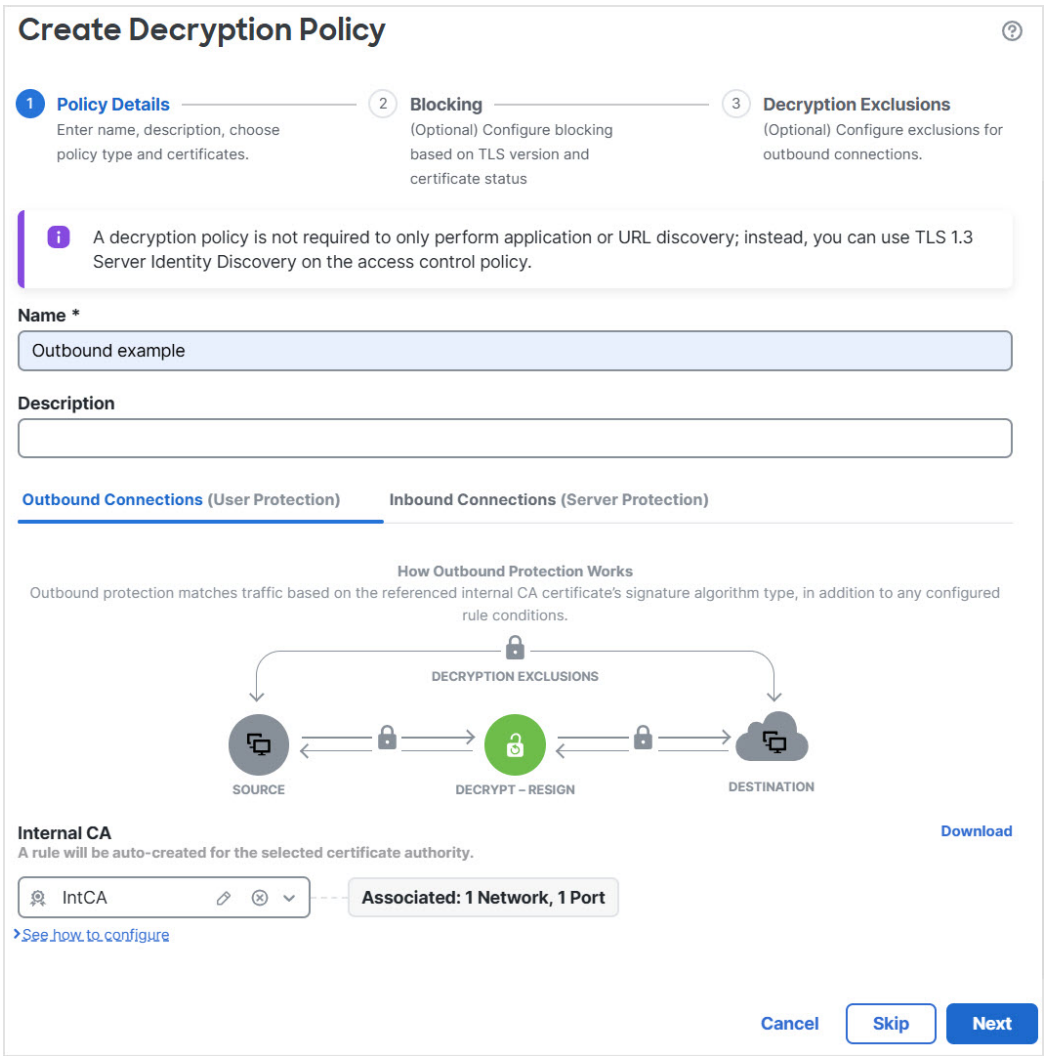

Step 7 From the **Internal CA** list, upload or choose certificates for the rules.

For more information about internal certificates, see Generate an Internal CA for Outbound Protection, on page 2291 and Upload an Internal CA for Outbound Protection, on page 2292.

Step 8 (Optional.) Choose networks and ports.

For more information:

- Network Rule Conditions, on page 939
- Port Rule Conditions, on page 940

Step 9 Click **Next**.

Step 10 Continue with Decryption Policy Block Connections, on page 2287.

What to do next

- Add rule conditions: Decryption Rule Conditions, on page 2316
- Add a default policy action: Decryption Policy Default Actions, on page 2295
- Configure logging options for the default action as described in *Logging Connections with a Policy Default Action* in the Cisco Secure Firewall Management Center Administration Guide .
- Set advanced policy properties: Decryption Policy Advanced Options, on page 2298.
- Associate the decryption policy with an access control policy as described in Associating Other Policies with Access Control, on page 1762.
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Create a Create a Decryption Policy with Inbound Connection Protection

This task discusses how to create a decryption policy with a rule that protects inbound connections; that is, the destination server is inside your protected network. This type of rule has a **Decrypt - Known Key** rule action.

When you create a decryption policy, you can create multiple rules at the same time, including multiple **Decrypt - Known Key** rules and multiple **Decrypt - Resign** rules.

Before you begin

You can optionally upload an internal certificate for your internal server before you can create a decryption policy that protects inbound connections. You can do this in any of the following ways:

- Create an internal certificate object by going to **Objects** > **Object Management** > **PKI** > **Internal Certs** and referring to PKI, on page 1404.
- At the time you create this decryption policy.

If you enabled Change Management, you must create and assign a ticket before you can create a decryption policy. Before the decryption policy can be used, the ticket and all associated objects (like certificate authorities) must be approved. For more information, see Creating Change Management Tickets, on page 240 and Policies and Objects that Support Change Management, on page 235.

Procedure

Step 1 Log in to Secure Firewall Management Center if you haven't already done so.

Step 2 Click **Policies** > **Access Control heading** > **Decryption**.

Step 3 Click **Create Decryption Policy**.

Step 4 Give the policy a unique **Name** and, optionally, a **Description**.

The following characters are not supported in decryption policy names:

- Leading period
- `#,;,{,},=,$,<,>`

Step 5 From the **Internal Certificates** list, upload or choose certificates for the rules.

For more information about internal CA certificates, see Internal Certificate Authority Objects, on page 1405.

Step 6 (Optional.) Choose networks and ports.

For more information:

- Network Rule Conditions, on page 939
- Port Rule Conditions, on page 940

Step 7 Click the **Inbound Connections** tab.

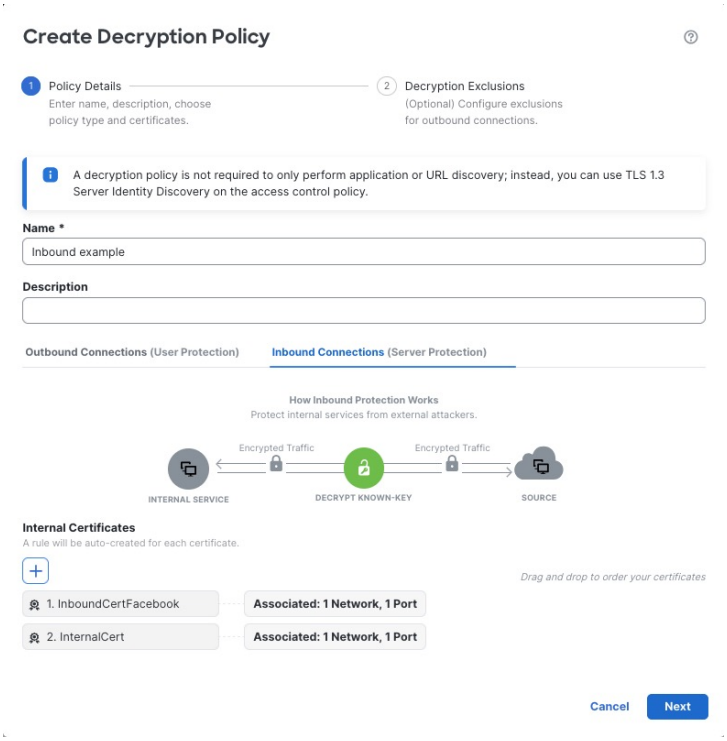

Step 8 Click **Next**.

Step 9 Continue with Decryption Policy Block Connections, on page 2287

Step 10 Continue with Decryption Policy Exclusions, on page 2287.

What to do next

- Add rule conditions: Decryption Rule Conditions, on page 2316
- Add a default policy action: Decryption Policy Default Actions, on page 2295
- Configure logging options for the default action as described in *Logging Connections with a Policy Default Action* in the Cisco Secure Firewall Management Center Administration Guide .
- Set advanced policy properties: Decryption Policy Advanced Options, on page 2298.

- Associate the decryption policy with an access control policy as described in Associating Other Policies with Access Control, on page 1762.
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Decryption Policy Block Connections

This topic provides details about how to block connections to servers with unsecure TLS versions and server certificate statuses while creating a decryption policy. The **Block with Reset** rules are created in your decryption policy that are disabled by default.

Procedure

Step 1 Complete the tasks mentioned in:

- Create a Create a Decryption Policy with Outbound Connection Protection, on page 2281
- Create a Create a Decryption Policy with Inbound Connection Protection, on page 2285

Step 2 The **Blocking** page provides the following options. By default, all the options are *disabled* for decryption policy actions.

- **Block connections based on TLS version**—Check this check box to block connections to servers using unsecure TLS versions. By default, **SSL v3.0**, **TLS v1.0**, and **TLS v1.1** which are known to be vulnerable, are selected. You can choose other versions from the drop-down list.
- **Block connections based on server certificate status**—Check this check box to block connections to servers with unsecure server certificate statuses. By default, **Invalid Signature**, **Expired**, **Not Yet Valid**, and **Invalid Certificate** are selected. You can choose other statuses from the drop-down list.

Click **Delete** ( ) to remove the selections or click **Reset to default** to revert back to the default selections.

Step 3 Click **Next**.

What to do next

Continue with Decryption Policy Exclusions, on page 2287.

Decryption Policy Exclusions

This task discusses how to exclude certain types of traffic from decryption. We create **Do Not Decrypt** rules in your decryption policy for these although the rules are initially enabled only for an outbound decryption policy (that is, one that uses the **Decrypt - Resign** policy action).

Before you begin

You must upload an internal CA certificate for your managed device before you can create a decryption policy that protects outbound connections. You can do this in any of the following ways:

- Create an internal CA certificate object by going to **Objects** > **Object Management** > **PKI** > **Internal CAs** and referring to PKI, on page 1404.
- At the time you create this decryption policy.

Procedure

Step 1 Complete the tasks discussed in:

- Create a Create a Decryption Policy with Outbound Connection Protection, on page 2281
- For more information, see Create a Create a Decryption Policy with Inbound Connection Protection, on page 2285

Step 2 The exclusions page provides the following options. All options are *enabled* for an outbound protection policy (**Decrypt - Resign** rule action) and *disabled* for all other decryption policy actions.

Item	Description
Bypass decryption for sensitive URL categories	Check the box to not decrypt traffic from the indicated categories. Depending on the laws in your area, decryption certain traffic, such as finance or health-related, might be prohibited. Consult an authority in your area for more information. Click **Add** to add more categories. Click **Delete** (✕) to remove categories.
Bypass decryption for undecryptable distinguished names	Check the box to not decrypt traffic when re-signing the certificate is likely to cause the connection to fail. Typically, this behavior is associated with *certificate pinning*, which is discussed in TLS/SSL Certificate Pinning Guidelines, on page 2310. The list of undecryptable distinguished names is maintained by Cisco.
Bypass decryption for undecryptable applications	Check the box to not decrypt traffic when re-signing the certificate is likely to cause the connection to fail. Typically, this behavior is associated with *certificate pinning*, which is discussed in TLS/SSL Certificate Pinning Guidelines, on page 2310. Undecryptable applications are updated automatically in the Vulnerability Database (VDB). You can find a list of all applications on the Secure Firewall Application Detectors page; the **undecryptable** tag identifies applications Cisco determines are undecryptable. The list of undecryptable applications is maintained by Cisco.

Item	Description
Bypass decryption for very low-risk connections	Check this check box to not decrypt traffic for very low-risk clients connecting to trusted servers, based on the threat confidence levels of clients which is determined by the Encrypted Visibility Engine (EVE) and the URL Category Reputation. An **Auto-Rule-Low-Risk-Connections** Do Not Decrypt rule is created and enabled in the new decryption policy. To use the **Bypass decryption for very low-risk connections** option, you must enable the EVE in the corresponding access control policy, and the managed devices must be running Version 7.7 or later with a valid IPS license.

The following figure shows the default options.

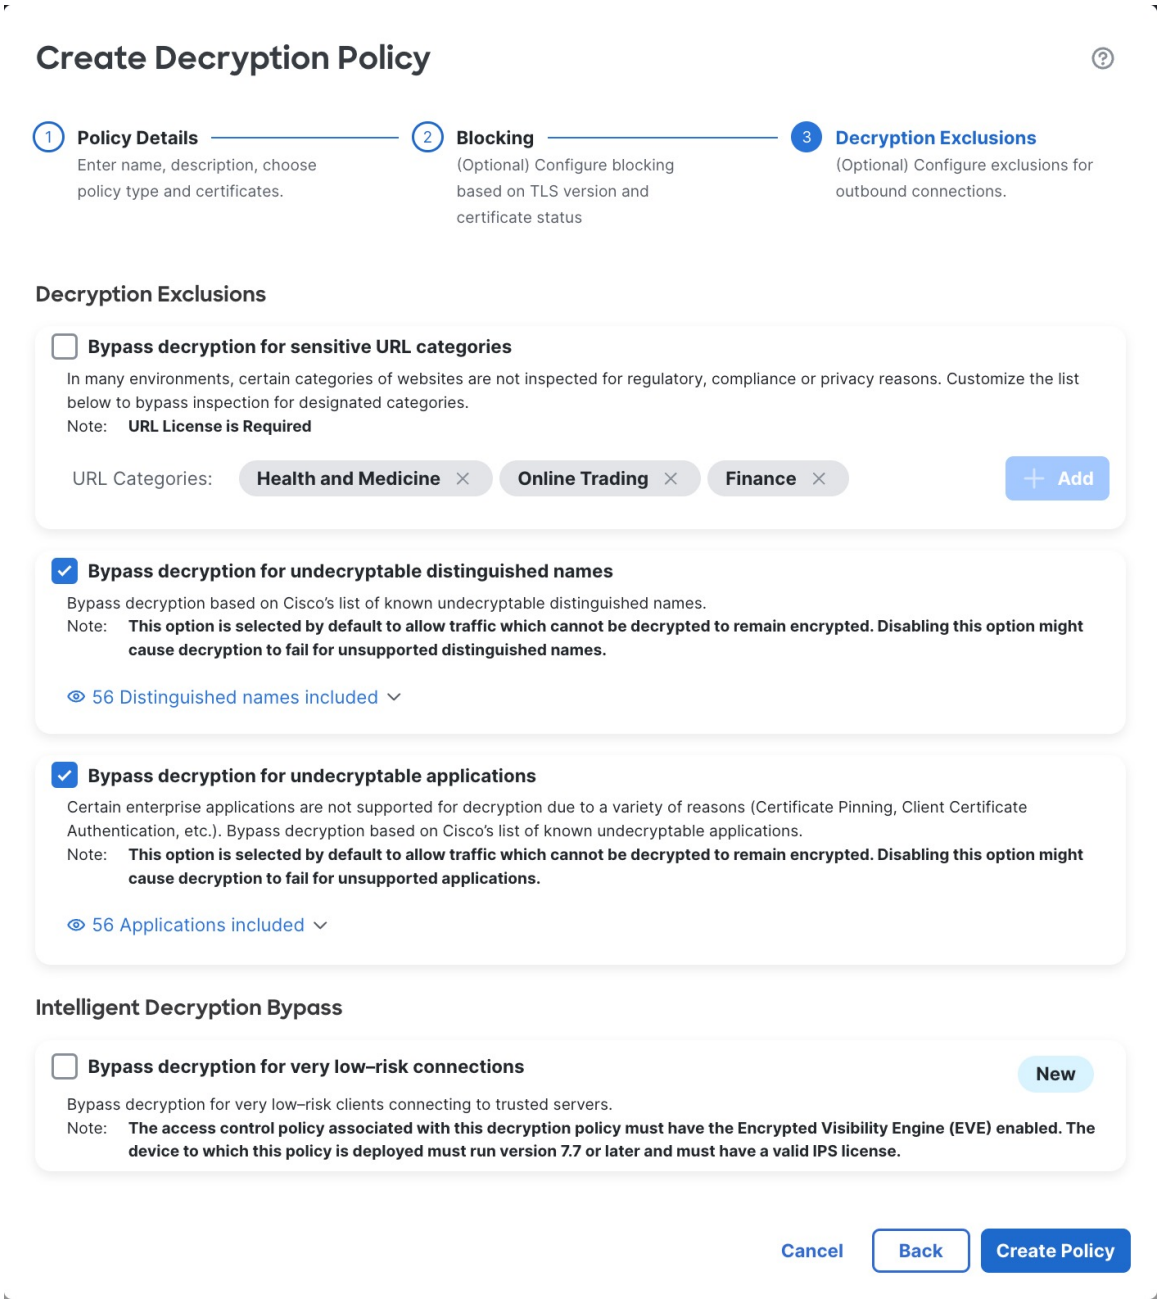

Step 3 Click **Create Policy**.

The following figure shows a sample outbound protection policy.

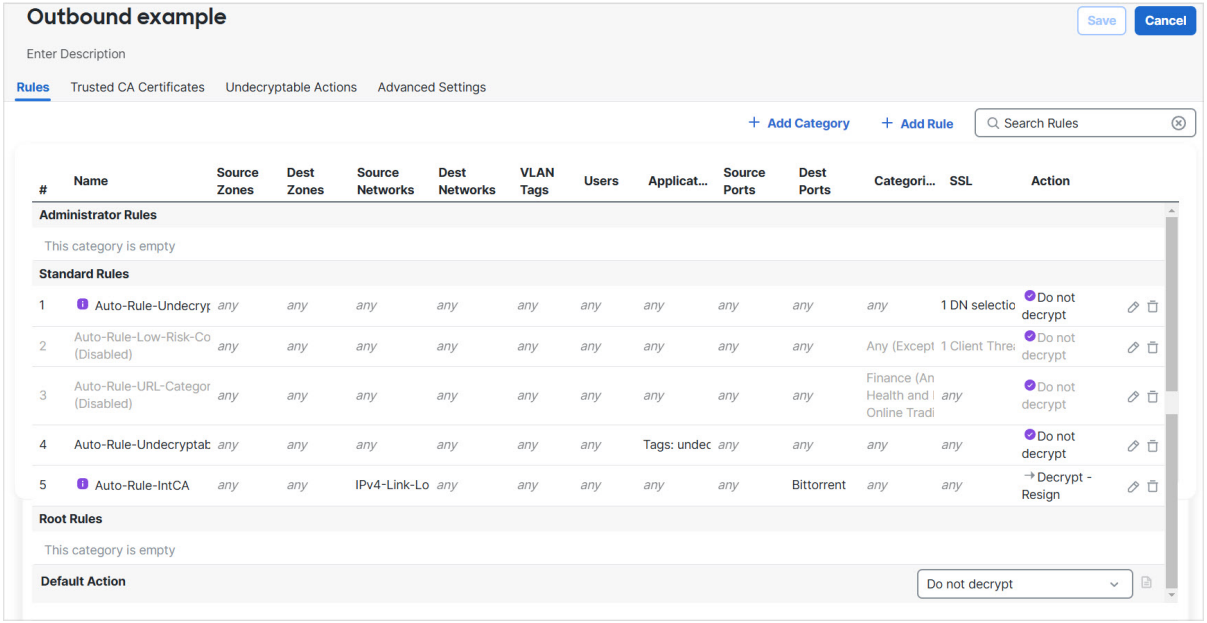

In the preceding example, the **Do Not Decrypt** rules corresponding to your choices for rule exclusions are automatically added before the **Decrypt - Resign** rule. The rule for sensitive URL categories is disabled because, by default, that exclusion is disabled. Had you selected the **Bypass decryption for sensitive URL categories** check box, the rule would have been enabled.

Step 4 Click **Create Policy**.

What to do next

- Add rule conditions: Decryption Rule Conditions, on page 2316
- Add a default policy action: Decryption Policy Default Actions, on page 2295
- Configure logging options for the default action as described in *Logging Connections with a Policy Default Action* in the Cisco Secure Firewall Management Center Administration Guide .
- Set advanced policy properties: Decryption Policy Advanced Options, on page 2298.
- Associate the decryption policy with an access control policy as described in Associating Other Policies with Access Control, on page 1762.
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Generate an Internal CA for Outbound Protection

This task discusses how you can optionally generate an internal certificate authority when you create a decryption rule that protects outbound connections. You can also perform these tasks using **Objects** > **Object Management** as discussed in Uploading a Signed Certificate Issued in Response to a CSR, on page 1407.

Before you begin

Make sure you understand the requirements for generating an internal certificate authority object as discussed in Internal Certificate Authority Objects, on page 1405.

Procedure

Step 1 Log in to Secure Firewall Management Center if you haven't already done so.

Step 2 Click **Policies** > **Access Control heading** > **Decryption**.

Step 3 Click **Create Decryption Policy (Rule-Based)**.

Step 4 Enter a name for the policy in the **Name** field and an optional description in the **Description** field.

Step 5 Click the **Outbound Connections** tab.

Step 6 From the **Internal CA** list, click **Create New** > **Generate CA**.

Step 7 Give the internal CA a **Name** and provide a two-letter **Country Name**.

Step 8 Click **Self-Signed** or **CSR**.

For more information about these options, see Internal Certificate Authority Objects, on page 1405.

Step 9 Enter the requested information in the provided fields.

Step 10 Click **Save**.

Step 11 If you chose **CSR**, after the signing request has been completed, click **Install Certificate** as follows:

a) Repeat the preceding steps in this procedure.

b) Edit the CA from the **Internal CA** list as follows.

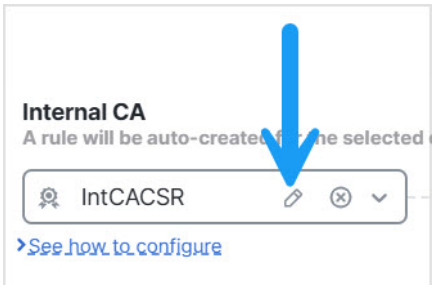

c) Click **Install Certificate**.

d) Follow the prompts on your screen to complete the task.

Step 12 Continue creating the policy as discussed in Create a Create a Decryption Policy with Outbound Connection Protection, on page 2281.

Upload an Internal CA for Outbound Protection

This task discusses how you can optionally upload an internal certificate authority when you create a decryption rule that protects outbound connections. You can also perform these tasks using **Objects** > **Object Management** as discussed in Uploading a Signed Certificate Issued in Response to a CSR, on page 1407.

Before you begin

Make sure you understand the requirements for generating an internal certificate authority object as discussed in Internal Certificate Authority Objects, on page 1405.

Procedure

Step 1 Log in to Secure Firewall Management Center if you haven't already done so.

Step 2 Click **Policies** > **Access Control heading** > **Decryption**.

Step 3 Click **Create Decryption Policy**.

Step 4 Enter a name for the policy in the **Name** field and an optional description in the **Description** field.

Step 5 Click the **Outbound Connections** tab.

Step 6 From the **Internal CA** list, click **Create New** > **Upload CA**.

Step 7 Give the internal CA a **Name**.

Step 8 Paste or browse to locate the certificate and its private key in the provided fields.

Step 9 If the CA has a password, select the **Encrypted** check box and enter the password in the adjacent field.

Step 10 Continue creating the policy as discussed in Create a Create a Decryption Policy with Outbound Connection Protection, on page 2281.

Upload an Internal Certificate for Inbound Protection

This task discusses how to upload an internal certificate authority when you create a decryption rule that protects outbound connections. You can also upload the internal CA using **Objects** > **Object Management** as discussed in Importing a CA Certificate and Private Key, on page 1406.

Before you begin

Make sure you have an internal certificate authority in one of the formats discussed in Internal Certificate Authority Objects, on page 1405.

Procedure

Step 1 Log in to Secure Firewall Management Center if you haven't already done so.

Step 2 Click **Policies** > **Access Control heading** > **Decryption**.

Step 3 Click **Create Decryption Policy**.

Step 4 Enter a name for the policy in the **Name** field and an optional description in the **Description** field.

Step 5 Click the **Inbound Connections** tab.

Step 6 From the **Internal Certificates** list, click **Add** (+).

Step 7 If an internal certificate object exists, click its name.

Step 8 Otherwise, click **Upload**.

Step 9 Enter the required information.

See Adding Internal Certificate Objects, on page 1412.

Step 10 Continue creating the decryption policy as discussed in Create a Create a Decryption Policy with Inbound Connection Protection, on page 2285.

Create a Create a Decryption Policy with Other Rule Actions

To create a decryption rule with a **Do Not Decrypt**, **Block**, **Block With Reset,** or **Monitor** rule action, create a decryption policy and edit the policy to add the rule.

When you create a decryption policy policy, you can create multiple rules at the same time, including multiple **Decrypt - Known Key** rules, multiple **Decrypt - Replace Cert** rules, and multiple **Decrypt - Resign** rules.

If you enabled Change Management, you must create and assign a ticket before you can create a decryption policy. Before the decryption policy can be used, the ticket and all associated objects (like certificate authorities) must be approved. For more information, see Creating Change Management Tickets, on page 240 and Policies and Objects that Support Change Management, on page 235.

Procedure

Step 1 Log in to Secure Firewall Management Center if you haven't already done so.

Step 2 Click **Policies** > **Access Control heading** > **Decryption**.

Step 3 Give the policy a unique **Name** and, optionally, a **Description**.

The following characters are not supported in decryption policy names:

- Leading period
- `#,;,{,},=,$,<,>`

Step 4 Click **Next**.

Step 5 The bypass page is provided for your information only; you cannot bypass traffic for other types of decryption (such as **Block**).

Step 6 Click **Create Policy**.

Step 7 Wait for the policy to be created.

Step 8 Click **Edit** (✎) next to the decryption policy name.

Step 9 Click **Add Rule**.

Step 10 Give the rule a **Name**.

Step 11 From the **Action** list, click a rule action and see one of the following sections for more information:

- Decryption Rule Do Not Decrypt Action, on page 2335
- Decryption Rule Blocking Actions, on page 2336
- Decryption Rule Monitor Action, on page 2334

Step 12 Click **Save**.

What to do next

- Add rule conditions: Decryption Rule Conditions, on page 2316
- Add a default policy action: Decryption Policy Default Actions, on page 2295
- Configure logging options for the default action as described in *Logging Connections with a Policy Default Action* in the Cisco Secure Firewall Management Center Administration Guide .
- Set advanced policy properties: Decryption Policy Advanced Options, on page 2298.
- Associate the decryption policy with an access control policy as described in Associating Other Policies with Access Control, on page 1762.
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Decryption Policy Default Actions

The default action for a decryption policy determines how the system handles decryptable encrypted traffic that does not match any non-monitor rule in the policy. When you deploy a decryption policy that does not contain any decryption rules, the default action determines how all decryptable traffic on your network is handled. Note that the system does not perform any kind of inspection on encrypted traffic blocked by the default action.

To set the decryption policy default action:

1. Log in to the Secure Firewall Management Center if you haven't already done so.
2. Click **Policies** > **Access Control heading** > **Decryption**.
3. Click **Edit** (✎) next to the name of the decryption policy.
4. In the Default Action row, click one of the following actions from the list.

Table 124: Decryption Policy Default Actions

Default Action	Effect on Encrypted Traffic
Block	Block the TLS/SSL session without further inspection.
Block with reset	Block the TLS/SSL session without further inspection and reset the TCP connection. Choose this option if traffic uses a connectionless protocol like UDP. In that case, the connectionless protocol tries to reestablish the connection until it is reset. This action also displays a connection reset error in the browser so the user is informed that the connection is blocked.
Do not decrypt	Inspect the encrypted traffic with access control.

Default Handling Options for Undecryptable Traffic

Table 125: Undecryptable Traffic Types

Type	Description	Default Action	Available Action
Compressed Session	The TLS/SSL session applies a data compression method.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
SSLv2 Session	The session is encrypted with SSL version 2. Note that traffic is decryptable if the ClientHello message is SSL 2.0, and the remainder of the transmitted traffic is SSL 3.0.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Unknown Cipher Suite	The system does not recognize the cipher suite.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Unsupported Cipher Suite	The system does not support decryption based on the detected cipher suite.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Session not cached	The TLS/SSL session has session reuse enabled, the client and server reestablished the session with the session identifier, and the system did not cache that session identifier.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Handshake Errors	An error occurred during TLS/SSL handshake negotiation.	Inherit default action	Do not decrypt Block Block with reset Inherit default action
Decryption Errors	An error occurred during traffic decryption.	Block	Block Block with Reset

When you first create a decryption policy, logging connections that are handled by the default action is disabled by default. Because the logging settings for the default action also apply to undecryptable traffic handling, logging connections handled by the undecryptable traffic actions is disabled by default.

Note that if your browser uses certificate pinning to verify a server certificate, you cannot decrypt this traffic by re-signing the server certificate. For more information, see Decryption Rule Guidelines and Limitations, on page 2304.

If a decryption policy is associated with an access control policy that uses TCP state bypass, matching traffic is acted on based on the policy's configured **Undecryptable Actions** for **Handshake Errors**.

For example, if a decryption policy's **Handshake Errors** is set to **Block**, traffic matching the rule is blocked and the connection event's action is reported as a handshake error.

For more information about TCP state bypass, see:

- Configure TCP State Bypass, on page 1945
- Bypass TCP State Checks for Asymetrical Routing (TCP State Bypass), on page 1943

Related Topics

Set Default Handling for Undecryptable Traffic, on page 2297

Set Default Handling for Undecryptable Traffic

You can set undecryptable traffic actions at the decryption policy level to handle certain types of encrypted traffic the system cannot decrypt or inspect. When you deploy a decryption policy that contains no decryption rules, the undecryptable traffic actions determine how all undecryptable encrypted traffic on your network is handled.

Depending on the type of undecryptable traffic, you can choose to:

- Block the connection.
- Block the connection, then reset it. This option is preferrable for connectionless protocols like UDP, which keep trying to connect until the connection is blocked.
- Inspect the encrypted traffic with access control.
- Inherit the default action from the decryption policy.

Procedure

Step 1 Log in to Secure Firewall Management Center if you haven't already done so.

Step 2 Click **Policies** > **Access Control heading** > **Decryption**.

Step 3 Click **Edit** (✎) next to the name of the decryption policy.

Step 4 In the decryption policy editor, click **Undecryptable Actions**.

Step 5 For each field, choose either the decryption policy's default action or another action you want to take on the type of undecryptable traffic. See Default Handling Options for Undecryptable Traffic, on page 2296 and Decryption Policy Default Actions, on page 2295 for more information.

Step 6 Click **Save** to save the policy.

What to do next

- Configure default logging for connections handled by the undecryptable traffic actions; see *Logging Connections with a Policy Default Action* in the Cisco Secure Firewall Management Center Administration Guide.
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Decryption Policy Advanced Options

A decryption policy 's **Advanced Settings** page has global settings that are applied to all managed devices that are configured for Snort 3 to which the policy is applied.

A decryption policy advanced settings are all ignored on any managed device that runs:

- A version earlier than 7.1
- Snort 2

Block flows requesting ESNI

Encrypted Server Name Indication (ESNI (link to draft proposal)) is a way for a client to tell a TLS 1.3 server what the client is requesting. Because the SNI is encrypted, you can optionally block these connections because the system cannot determine what the server is.

Disable HTTP/3 advertisement

This option strips HTTP/3 (RFC 9114) from the ClientHello in TCP connections. HTTP/3 is part of the QUIC transport protocol, not the TCP transport protocol. Blocking clients from advertising HTTP/3 provides protection against attacks and evasion attempts potentially burried within QUIC connections.

Propagate untrusted server certificates to clients

This applies only to traffic matching a **Decrypt - Resign** rule action.

Enable this option to substitute the certificate authority (CA) on the managed device for the server's certificate in cases where the server certificate is untrusted. An *untrusted* server certificate is one that is not listed as a trusted CA in the Secure Firewall Management Center. (**Objects** > **Object Management** > **PKI** > **Trusted CAs**).

Enable TLS 1.3 Decryption

Whether to apply decryption rules to TLS 1.3 connections. If you do not enable this option, the decryption rules apply to TLS 1.2 or lower traffic only. See TLS 1.3 Decryption Best Practices, on page 2300.

Enable adaptive TLS server identity probe

Automatically enabled when TLS 1.3 decryption is enabled. A *probe* is a partial TLS connection with the server, the purpose of which is to obtain the server certificate and cache it. (If the certificate is already cached, the probe is never established.)

If TLS 1.3 Server Identity Discovery is disabled on the access control policy with which the decryption policy is associated, we attempt to use the Server Name Indication (SNI), which is not as reliable.

The adaptive TLS server identity probe occurs on any of the following conditions as opposed to on every connection as in earlier releases:

- Certificate Issuer—Matched when the value of **Issuer DNs** in a decryption rule's DN rule condition is matched.

  For more information, see Distinguished Name (DN) Rule Conditions, on page 2323.

- Certificate Status—Matched when any of the **Cert Status** conditions are matched in a decryption rule.

  For more information, see Certificate Status Decryption Rule Conditions, on page 2329.

- Internal/External Certificate—Internal certificates can be matched by the certificate used in **Decrypt - Known Key** rule actions; external certificates can be matched in **Certificates** rule conditions.

  For more information, see Known Key Decryption (Incoming Traffic), on page 2263 and Certificate Decryption Rule Conditions, on page 2322.

- Application ID—Can be matched by **Applications** rule conditions in either an access control policy or a decryption policy.

  For more information, see Application Rule Conditions, on page 939.

- URL Category—Can be matched by **URLs** rule conditions in an access control policy.

  For more information, see URL Rule Conditions, on page 942.

Note **Enable adaptive TLS server discovery mode** is not supported on any Secure Firewall Threat Defense Virtual deployed to AWS. If you have any such managed devices managed by the Secure Firewall Management Center, the connection event **PROBE_FLOW_DROP_BYPASS_PROXY** increments every time the device attempts to extract the server certificate.

Enable QUIC Decryption

Whether to apply decryption rules to connections that use the HTTP/3 over the QUIC protocol. When you decrypt QUIC connections, the system can inspect the contents of the sessions for intrusions, malware, or other issues. You can also apply granular control and filtering of decrypted QUIC connections based on specific criteria in the access control policy. QUIC support is in line with RFC 9000, 9001, 9002, 9114, 9204.

Consider the following when implementing QUIC decryption:

- On high availability or clustered devices, QUIC decryption works only if the connection remains on the same node. If the connection fails over, or is forwarded to another node, the connection drops and must be re-established. Multi-instance is supported without restrictions.
- Rules that apply to QUIC traffic would include the UDP protocol with destination port 443.

- Access control rules that apply to QUIC traffic would include the HTTP/3 or QUIC protocols, either explicitly or by implication.

The following limitations apply to QUIC decryption:

- QUIC decryption applies to Firewall Threat Defense 7.6+ only. Devices running a lower release cannot decrypt QUIC connections.
- Connections from browsers using the Chromium stack (Google Chrome, Opera, Edge) cannot be decrypted for outbound traffic. But inbound traffic from the same browsers can be decrypted.
- Connection Migration as described in RFC 9000 is not supported. The concept of Connection ID in QUIC allows endpoints to retain the same connection in the event of address change.
- Key update, session resumption, and QUIC version 2 are not supported.
- Interactive Block and Interactive Block with Reset (in access control rules) is not supported. These actions will work as Block and Block with Reset.
- The active connection-ID per connection is limited to 5. If necessary, you can modify these limits using the **system support quic-tuning** and **system support quic-tuning-reset** commands in the device CLI.

TLS 1.3 Decryption Best Practices

Recommendation: When to enable advanced options

Both the decryption policy and the access control policy have advanced options that affect how traffic is handled, whether the traffic is being decrypted or not.

The advanced options are:

- Decryption policy:
  - TLS 1.3 decryption
  - TLS adaptive server identity probe
- Access control policy: TLS 1.3 Server Identity Discovery

  The access control policy setting takes precedence over the decryption policy setting.

Use the following table to decide which option to enable:

TLS adaptive server identity probe setting (decryption policy)	TLS 1.3 Server Identity Discovery setting (access control policy)	Result	Recommended when
Enabled	Disabled	Adaptive probe sent if decryption policy contains *any* rule conditions specified in Decryption Policy Advanced Options, on page 2298 *and* if the server certificate is not cached.	• You're not using application or URL conditions in access control rules • You're decrypting traffic

TLS adaptive server identity probe setting (decryption policy)	TLS 1.3 Server Identity Discovery setting (access control policy)	Result	Recommended when
Enabled	Enabled	Probe is always sent if the server certificate is not cached.	Use only if your access control rules have URL or application conditions
Disabled	Enabled	Probe is always sent if the server certificate is not cached.	Not recommended.
Disabled	Disabled	Probe is never sent.	Very limited usefulness; use only if not decrypting traffic and not using application or URL conditions in the access control rule

Note A cached TLS server's certificate is available to all Snort instances on a particular Firewall Threat Defense. The cache can be cleared with a CLI command and is automatically cleared when the device is rebooted.

Reference

For more information, see the discussion of TLS server identity discovery on secure.cisco.com.

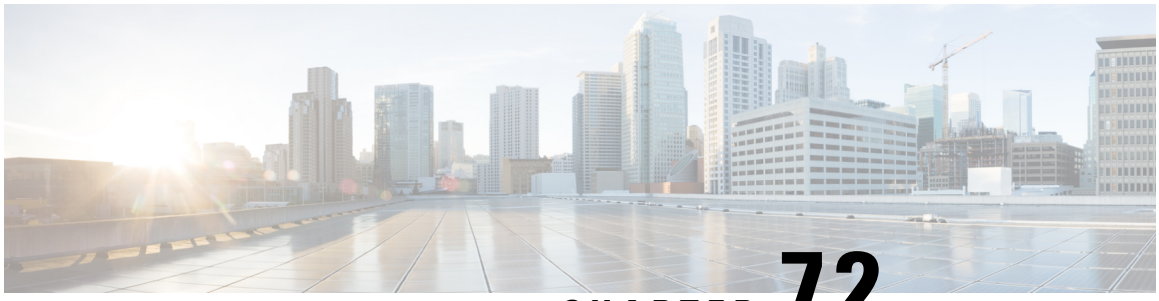

CHAPTER 72

Decryption Rules

The following topics provide an overview of creating, configuring, managing, and troubleshooting decryption rules:

Note Because TLS and SSL are often used interchangeably, we use the expression *TLS/SSL* to indicate that either protocol is being discussed. The SSL protocol has been deprecated by the IETF in favor of the more secure TLS protocol, so you can usually interpret *TLS/SSL* as referring to TLS only.

For more information about SSL and TLS protocols, see a resource such as SSL vs. TLS - What's the Difference?

- Decryption Rules Overview, on page 2303
- Requirements and Prerequisites for Decryption Rules, on page 2303
- Decryption Rule Guidelines and Limitations, on page 2304
- Decryption Rule Traffic Handling, on page 2312
- Decryption Rule Conditions, on page 2316
- Decryption Rule Actions, on page 2334
- Monitor TLS/SSL Hardware Acceleration, on page 2337
- Troubleshoot Decryption Rules, on page 2339
- Troubleshooting Using Crypto Archives, on page 2349

Decryption Rules Overview

Decryption rules provide a granular method of handling encrypted traffic across multiple managed devices, whether blocking the traffic without further inspection, not decrypting the traffic and inspecting it with access control, or decrypting the traffic for access control analysis.

Requirements and Prerequisites for Decryption Rules

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Decryption Rule Guidelines and Limitations

Keep the following points in mind when setting up your decryption rules. Properly configuring decryption rules is a complex task, but one that is essential to building an effective deployment that handles encrypted traffic. Many factors influence how you configure rules, including certain application behavior that you cannot control.

In addition, rules can preempt each other, require additional licenses, or contain invalid configurations. Thoughtfully configured rules can also reduce the resources required to process network traffic. Creating overly complex rules and ordering rules the wrong way can adversely affect performance.

For detailed information, see Best Practices for Access Control Rules, on page 1738.

For guidelines related specifically to TLS crypto acceleration, see TLS Crypto Acceleration, on page 2268.

Related Topics

Rule and Other Policy Warnings
Best Practices for Access Control Rules, on page 1738
Guidelines for Using TLS/SSL Decryption, on page 2304
Decryption Rule Unsupported Features, on page 2305
TLS/SSL Do Not Decrypt Guidelines, on page 2306
TLS/SSL Decrypt - Resign Guidelines, on page 2307
TLS/SSL Decrypt - Known Key Guidelines, on page 2309
TLS/SSL Block Guidelines, on page 2310
TLS/SSL Certificate Pinning Guidelines, on page 2310
TLS/SSL Heartbeat Guidelines, on page 2311
TLS/SSL Anonymous Cipher Suite Limitation, on page 2311
TLS/SSL Normalizer Guidelines, on page 2311
Other Decryption Rule Guidelines, on page 2311
Decryption Policy Rule Order

Guidelines for Using TLS/SSL Decryption

General guideline

decryption rules require processing overhead that can impact performance. Determine which traffic must be decrypted and subjected to deep inspection before you set up any policies or rules.

You cannot decrypt traffic on a device that uses any of the following:

- Passive interface
- Inline or inline tap interface

- TCP state bypass

Guidelines for undecryptable traffic

We can determine that certain traffic is not decryptable either because the website itself is not decryptable or because the website uses TLS/SSL pinning, which effectively prevents users from accessing a decrypted site without errors in their browser.

For more information about certificate pinning, see About TLS/SSL Pinning, on page 2343.

We maintain the list of these sites as follows:

- A Distinguished Name (DN) group named **Cisco-Undecryptable-Sites**
- The **pinned certificate** or **undecryptable** application filter

If you are decrypting traffic and you do not want users to see errors in their browsers when going to these sites, we recommend you set up a **Do Not Decrypt** rule toward the bottom of your decryption rules.

If you use the decryption policy wizard to create a policy for outbound traffic protection, a **Do Not Decrypt** rule for pinned certificates is created for you as the following example shows.

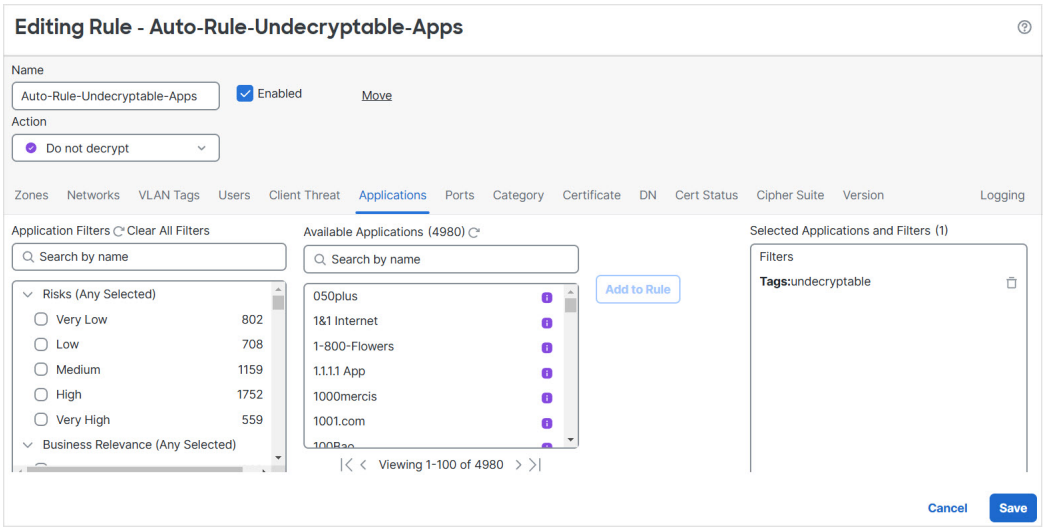

Decryption Rule Unsupported Features

RC4 cipher not unsupported

The Rivest Cipher 4 (also referred to as *RC4* or *ARC4* cipher is known to have vulnerabilities and is considered unsecure.

decryption policies identify the RC4 cipher suite as unsupported; you should configure the **Unsupported Cipher Suite** action in the decryption policies policy's **Undecryptable Actions** page to match your organization's requirements. For more information, see Default Handling Options for Undecryptable Traffic, on page 2296.

Passive, inline tap mode, and SPAN interfaces not supported

TLS/SSL traffic cannot be decrypted on passive, inline tap mode, or SPAN interfaces.

TLS/SSL Do Not Decrypt Guidelines

You should not decrypt traffic if doing so is forbidden by:

- Law; for example, some jurisdictions forbid decrypting financial information
- Company policy; for example, your company might forbid decrypting privileged communications
- Privacy regulations
- Traffic that uses certificate pinning (also referred to as *TLS/SSL pinning*) must remain encrypted to prevent breaking the connection

Encrypted traffic can be allowed or blocked on any decryption rule condition, including, but not limited to:

- Certificate status (for example, expired or invalid certificate)
- Protocol (for example, the nonsecure SSL protocol)
- Network (security zone, IP address, VLAN tag, and so on)
- URL category and reputation
- Port
- User group

Limitations of categories in Do Not Decrypt rules

You can optionally choose to include categories in your decryption policies. These categories, also referred to as *URL filtering*, are updated by the Cisco Talos intelligence group. Updates are based on machine learning and human analysis according to content that is retrievable from the website destination and sometimes from its hosting and registration information. Categorization is *not* based on the declared company vertical, intent, or security. While we strive to continuously update and improve URL filtering categories, it is not an exact science. Some websites are not categorized at all and it's possible some websites might be improperly categorized.

Avoid overusing categories in do not decrypt rules to avoid decrypting traffic without a reason; for example, the Health and Medicine category includes the WebMD website, which does not threaten patient privacy.

Following is a sample decryption policy that can prevent decryption for websites in the Health and Medicine categories but allows decryption for WebMD and everything else. General information about decryption rules can be found in Guidelines for Using TLS/SSL Decryption, on page 2304.

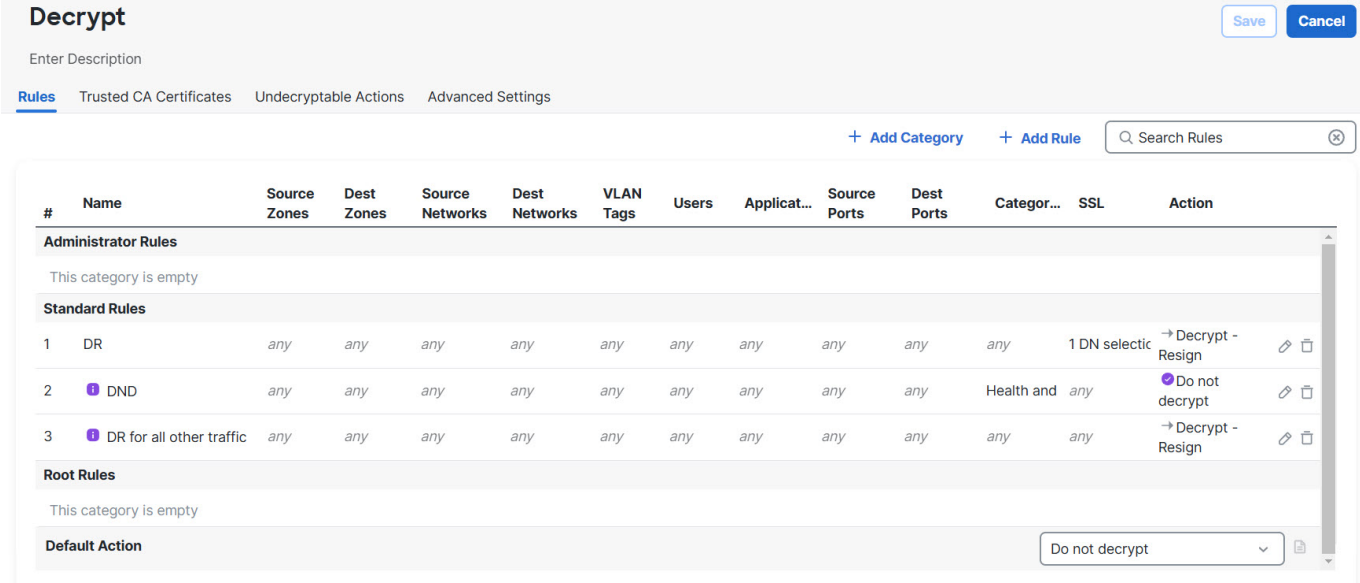

Note Don't confuse URL filtering with application detection, which relies on reading some of the packet from a website to determine more specifically what it is (for example, Facebook Message or Salesforce). For more information, see Best Practices for Configuring Application Control, on page 1735.

TLS/SSL Decrypt - Resign Guidelines

You can associate one internal Certificate Authority (CA) certificate and private key with the **Decrypt - Resign** action. If traffic matches this rule, the system re-signs the server certificate with the CA certificate, then acts as a man-in-the-middle. It creates two TLS/SSL sessions, one between client and managed device, one between managed device and server. Each session contains different cryptographic session details, and allows the system to decrypt and reencrypt traffic.

Best practices

We recommend the following:

- Use the **Decrypt - Resign** rule action for decrypting *outgoing* traffic, as opposed to incoming traffic for which we recommend the **Decrypt - Known Key** rule action.

  For more information about **Decrypt - Known Key** , see TLS/SSL Decrypt - Known Key Guidelines, on page 2309.

- Always check the **Replace Key Only** check box when you set up a **Decrypt - Resign** rule action.

  When a user browses to a web site that uses a *self-signed* certificate, the user sees a security warning in the web browser and is aware that they are communicating with an unsecure site.

  When a user browses to a web site that uses a trusted certificate, the user does not see a security warning.

Details

If you configure a rule with the **Decrypt - Resign** action, the rule matches traffic based on the referenced internal CA certificate's signature algorithm type, in addition to any configured rule conditions. Because you associate one CA certificate with a **Decrypt - Resign** action, you cannot create a decryption rule that decrypts multiple types of outgoing traffic encrypted with different signature algorithms. In addition, any external certificate objects and cipher suites you add to the rule must match the associated CA certificate encryption algorithm type.

For example, outgoing traffic encrypted with an elliptic curve (EC) algorithm matches a **Decrypt - Resign** rule only if the action references an EC-based CA certificate; you must add EC-based external certificates and cipher suites to the rule to create certificate and cipher suite rule conditions.

Similarly, a **Decrypt - Resign** rule that references an RSA-based CA certificate matches only outgoing traffic encrypted with an RSA algorithm; outgoing traffic encrypted with an EC algorithm does not match the rule, even if all other configured rule conditions match.

Guidelines and limitations

Also note the following:

Anonymous cipher suite unsupported

By nature, anonymous cipher suites are not used for authentication and do not use key exchanges. There are limited uses for anonymous cipher suites; for more information, see RFC 5246, appendix F.1.1.1. (Replaced for TLS 1.3 by RFC 8446 appendix C.5.)

You cannot use the **Decrypt - Resign** or **Decrypt - Known Key** action in the rule because anonymous cipher suites are not used for authentication.

Do not use Version or Cipher Suite rule conditions

Important *Never* use either **Cipher Suite** or **Version** rule conditions in a rule with a **Decrypt - Resign** or **Decrypt - Known Key** rule action. The use of these conditions in rules with other rule actions can interfere with the system's ClientHello processing, resulting in unpredictable performance.

Decrypt - Resign rule action and a Certificate Signing Request

To use a **Decrypt - Resign** rule action, you should create a Certificate Signing Request (CSR) and have it signed by a trusted certificate authority. (You can use the FMC to create a CSR: **Objects** > **Object Management** > **PKI** > **Internal Certs**.)

To be used in a **Decrypt - Resign** rule, your certificate authority (CA) must have at least one of the following extensions:

- **CA: TRUE**

  For more information, see the discussion of Basic Constraints in RFC3280, section 4.2.1.10.

- **KeyUsage=CertSign**

  For more information see RFC 5280, section 4.2.1.3.

To verify your CSR or CA has at least one of the preceding extensions, you can use the **openssl** command as discussed in a reference such as the openssl documentation.

This is necessary because for **Decrypt - Resign** inspection to work, the certificate that used in the decryption policy generates certificates on-the-fly and signs them so as to act as man-in-the middle and proxy all TLS/SSL connections.

Certificate pinning

If the customer's browser uses certificate pinning to verify a server certificate, you cannot decrypt this traffic by re-signing the server certificate. To allow this traffic, configure a decryption rule with the **Do Not Decrypt** action to match the server certificate common name or distinguished name.

The decryption policy wizard creates such rules for you; see Create a Create a Decryption Policy with Outbound Connection Protection, on page 2281 .

Non-matching cipher suite

The following error is displayed if you attempt to save a decryption rule with a cipher suite that does not match the certificate. To resolve the issue, see Verify TLS/SSL Cipher Suites, on page 2348.

```
Traffic cannot match this rule; none of your selected cipher suites contain a
signature algorithm that the resigning CA's signature algorithm
```

Untrusted Certificate Authority

If the client does not trust the Certificate Authority (CA) used to re-sign the server certificate, it warns the user that the certificate should not be trusted. To prevent this, import the CA certificate into the client trusted CA store. Alternatively, if your organization has a private PKI, you can issue an intermediate CA certificate signed by the root CA which is automatically trusted by all clients in the organization, then upload that CA certificate to the device.

HTTP proxy limitation

The system cannot decrypt traffic if an HTTP proxy is positioned between a client and your managed device, and the client and server establish a tunneled TLS/SSL connection using the CONNECT HTTP method. The **Handshake Errors** undecryptable action determines how the system handles this traffic.

Upload signed CA

If you create an internal CA object and choose to generate a certificate signing request (CSR), you cannot use this CA for a **Decrypt - Resign** action until you upload the signed certificate to the object.

Mismatched signature algorithm

If you configure a rule with the **Decrypt - Resign** action, and mismatch signature algorithm type for one or more external certificate objects or cipher suites, the policy editor displays an **Information**( ) next to the rule. If you mismatch signature algorithm type for all external certificate objects, or all cipher suites, the policy displays a warning icon **Warning** ( ) next to the rule, and you cannot deploy the access control policy associated with the decryption policy.

TLS/SSL Decrypt - Known Key Guidelines

When you configure the **Decrypt - Known Key** action, you can associate one or more server certificates and paired private keys with the action. If traffic matches the rule, and the certificate used to encrypt the traffic matches the certificate associated with the action, the system uses the appropriate private key to obtain the session encryption and decryption keys. Because you must have access to the private key, this action is best suited to decrypt traffic incoming to servers your organization controls.

Also note the following:

Anonymous cipher suite unsupported

By nature, anonymous cipher suites are not used for authentication and do not use key exchanges. There are limited uses for anonymous cipher suites; for more information, see RFC 5246, appendix F.1.1.1. (Replaced for TLS 1.3 by RFC 8446 appendix C.5.)

You cannot use the **Decrypt - Resign** or **Decrypt - Known Key** action in the rule because anonymous cipher suites are not used for authentication.

Cannot match on Distinguished Name or Certificate

You cannot match on **Distinguished Name** or **Certificate** rule conditions when creating a decryption rule with a **Decrypt - Known Key** action. The assumption is that if this rule matches traffic, the certificate, subject DN, and issuer DN already match the certificate associated with the rule.

Do not use Version or Cipher Suite rule conditions

Important *Never* use either **Cipher Suite** or **Version** rule conditions in a rule with a **Decrypt - Resign** or **Decrypt - Known Key** rule action. The use of these conditions in rules with other rule actions can interfere with the system's ClientHello processing, resulting in unpredictable performance.

Elliptic Curve Digital Signature Algorithm (*ECDSA*) certificate results in blocked traffic
(TLS 1.3 decryption enabled only.) If you use an ECDSA certificate with a **Decrypt - Known Key** rule action, matching traffic will be blocked. To avoid this, use a certificate with another type of certificate.

TLS/SSL Block Guidelines

If your decryption rule is associated with an access control policy that has a rule with an action of either **Interactive Block** or **Interactive Block with reset**, the system displays a customizable response page.

Provided you enabled logging in your rule, two connection events are displayed (in **Analysis** > **Connections**): One event for the interactive block and another event to indicate whether or not the user chose to continue to the site or not.

Related Topics

Configure HTTP Response Pages, on page 1873

TLS/SSL Certificate Pinning Guidelines

Some applications use a technique referred to as *TLS/SSL pinning* or *certificate pinning*, which embeds the fingerprint of the original server certificate in the application itself. As a result, if you configured a decryption rule with a **Decrypt - Resign** action, when the application receives a resigned certificate from a managed device, validation fails and the connection is aborted.

Because TLS/SSL pinning is used to avoid man-in-the-middle attacks, there is no way to prevent or work around it. We recommend adding a **Do Not Decrypt** rule before the **Decrypt - Resign** rule so pinning traffic is excluded from being decrypted. The decryption policy wizard does this for you.

- Create a **Do Not Decrypt** for those applications rule ordered before **Decrypt - Resign** rules.
- Instruct users to access the applications using a web browser.

For more information about rule ordering, see Decryption Policy Rule Order.

To determine whether applications are using TLS/SSL pinning, see Troubleshoot TLS/SSL Pinning, on page 2344.

TLS/SSL Heartbeat Guidelines

Some applications use the *TLS heartbeat* extension to the Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) protocols defined by RFC6520. TLS heartbeat provides a way to confirm the connection is still alive—either the client or server sends a specified number of bytes of data and requests the other party echo the response. If this is successful, encrypted data is sent.

You can configure a **Max Heartbeat Length** in a Network Analysis Policy (NAP) to determine how to handle TLS heartbeats. For more information, see The SSL Preprocessor.

For more information, see About TLS Heartbeat, on page 2341.

TLS/SSL Anonymous Cipher Suite Limitation

By nature, anonymous cipher suites are not used for authentication and do not use key exchanges. There are limited uses for anonymous cipher suites; for more information, see RFC 5246, appendix F.1.1.1. (Replaced for TLS 1.3 by RFC 8446 appendix C.5.)

You cannot use the **Decrypt - Resign** or **Decrypt - Known Key** action in the rule because anonymous cipher suites are not used for authentication.

You can add an anonymous cipher suite to the **Cipher Suite** condition in a decryption rule, but the system automatically strips anonymous cipher suites during ClientHello processing. For the system to use the rule, you must also configure your decryption rules in an order that prevents ClientHello processing. For more information, see Decryption Policy Rule Order.

TLS/SSL Normalizer Guidelines

If you enable the **Normalize Excess Payload** option in the inline normalization preprocessor, when the preprocessor normalizes decrypted traffic, it might drop a packet and replace it with a trimmed packet. This does not end the TLS/SSL session. If the traffic is allowed, the trimmed packet is encrypted as part of the TLS/SSL session.

Other Decryption Rule Guidelines

Users and groups

If you add a group or user to a rule, then change your realm settings to exclude that group or user, the rule has no effect. (The same applies to disabling the realm.) For more information about realms, see Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.

Categories in decryption rules

If your decryption policy has a **Decrypt - Resign** action but web sites are not being decrypted, check **Category** page on rules associated with that policy.

In some cases, a web site redirects to another site for authentication or other purposes and the redirected site might have a different URL categorization than the site you're trying to decrypt. For example, `gmail.com` (**Web based email** category) redirects to `accounts.gmail.com` (**Internet Portals** category) for authentication. Be sure to include all relevant categories in the SSL rule.

Note In order to fully process traffic based on URL category, you must also configure URL filtering. See the URL Filtering, on page 1857 chapter.

Query for URLs not in the local database

If you create a **Decrypt - Resign** rule and users browse to a web site whose category and reputation are not in the local database, data might not be decrypted. Some web sites are not categorized in the local database and, if not, data from those web sites is not decrypted by default.

You can control this behavior with the **Query Cisco cloud for unknown URLs** setting in **Integration** > **Other Integrations** > **Cloud Services**

Decryption Rule Traffic Handling

The system matches traffic to decryption rules in the order you specify. In most cases, the system handles encrypted traffic according to the *first* decryption rule where *all* the rule's conditions match the traffic. Conditions can be simple or complex; you can control traffic by security zone, network or geographical location, VLAN, port, application, requested URL, user, certificate, certificate distinguished name, certificate status, cipher suite, or encryption protocol version.

Each rule also has an *action*, which determines whether you monitor, block, or inspect matching encrypted or decrypted traffic with access control. Note that the system does *not* further inspect encrypted traffic it blocks. It does inspect encrypted and undecryptable traffic with access control. However, some access control rule conditions require unencrypted traffic, so encrypted traffic may match fewer rules. Also, by default, the system disables intrusion and file inspection of encrypted payloads.

The following scenario summarizes the ways that decryption rules handle traffic in an inline deployment.

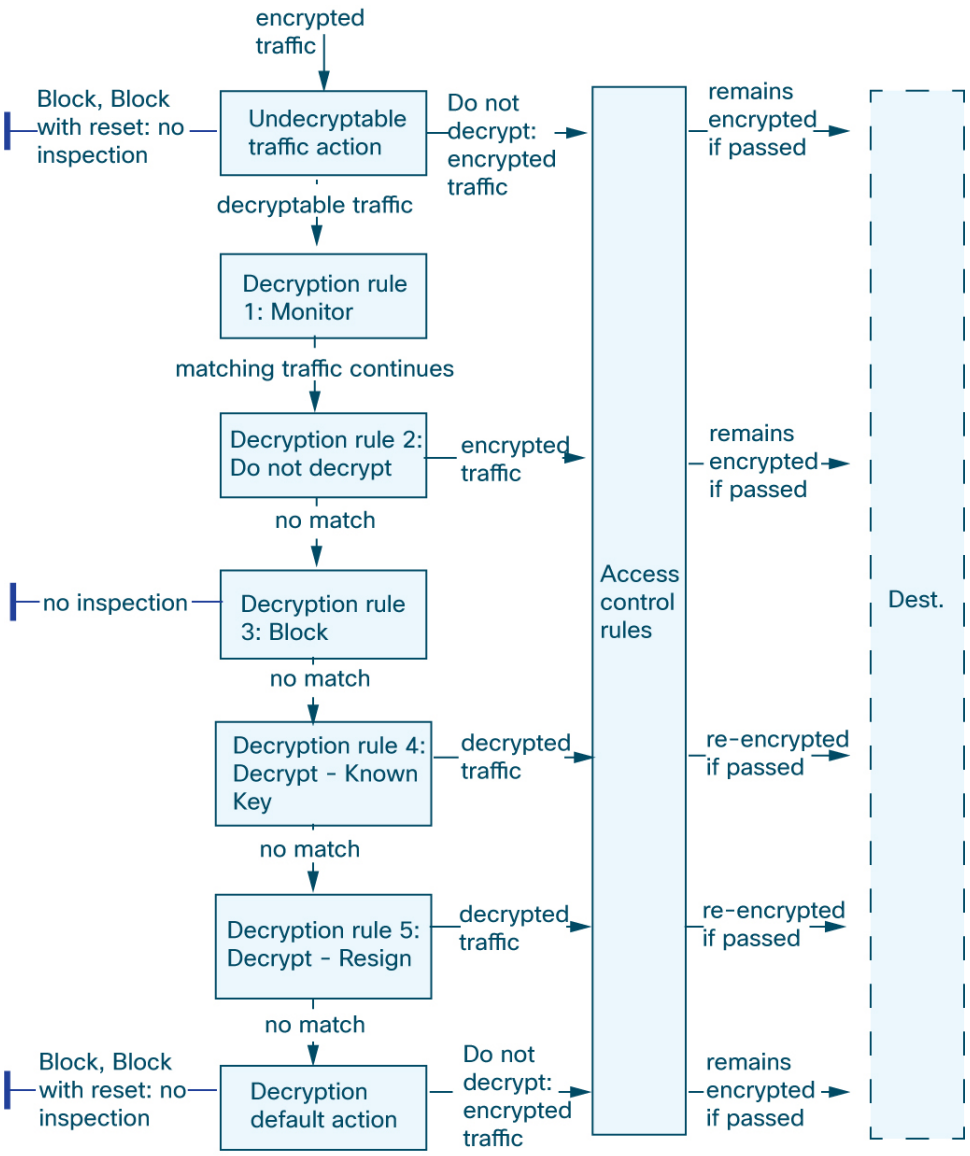

In this scenario, traffic is evaluated as follows:

- **Undecryptable Traffic Action** evaluates encrypted traffic first. For traffic the system cannot decrypt, the system either blocks it without further inspection or passes it for access control inspection. Encrypted traffic that does not match continues to the next rule.
- **Decryption Rule 1: Monitor** evaluates encrypted traffic next. Monitor rules track and log encrypted traffic but do not affect traffic flow. The system continues to match traffic against additional rules to determine whether to permit or deny it.
- **Decryption Rule 2: Do Not Decrypt** evaluates encrypted traffic third. Matching traffic is not decrypted; the system inspects this traffic with access control, but not file or intrusion inspection. Traffic that does not match continues to the next rule.

- **Decryption Rule 3: Block** evaluates encrypted traffic fourth. Matching traffic is blocked without further inspection. Traffic that does not match continues to the next rule.
- **Decryption Rule 4: Decrypt - Known Key** evaluates encrypted traffic fifth. Matching traffic incoming to your network is decrypted using a private key you upload. The decrypted traffic is then evaluated against access control rules. Access control rules handle decrypted and unencrypted traffic identically. The system can block traffic as a result of this additional inspection. All remaining traffic is reencrypted before being allowed to the destination. Traffic that does not match the decryption rule continues to the next rule.
- **Decryption Rule 5: Decrypt - Resign** is the final rule. If traffic matches this rule, the system re-signs the server certificate with an uploaded CA certificate, then acts as a man-in-the-middle to decrypt traffic. The decrypted traffic is then evaluated against access control rules. Access control rules treat decrypted and unencrypted traffic identically. The system can block traffic as a result of this additional inspection. All remaining traffic is reencrypted before being allowed to the destination. Traffic that does not match the decryption rule rule continues to the next rule.
- **Decryption policy Default Action** handles all traffic that does not match any of the decryption rules. The default action either blocks encrypted traffic without further inspection or does not decrypt it, passing it for access control inspection.

Encrypted Traffic Inspection Configuration

You must create reusable public key infrastructure (PKI) objects to control encrypted traffic based on encrypted session characteristics and decrypt encrypted traffic. You can add this information on the fly when uploading trusted certificate authority (CA) certificates to the a decryption policy and creating decryption rule, creating the associated object in the process. However, configuring these objects ahead of time reduces the chance of improper object creation.

Decrypting Encrypted Traffic with Certificates and Paired Keys

The system can decrypt incoming encrypted traffic if you configure an internal certificate object by uploading the server certificate and private key used to encrypt the session. If you reference that object in a decryption policy rule with an action of **Decrypt - Known Key** and traffic matches that rule, the system uses the uploaded private key to decrypt the session.

The system can also decrypt outgoing traffic if you configure an internal CA object by uploading a CA certificate and private key. If you reference that object in a decryption rule with an action of **Decrypt - Resign** and traffic matches that rule, the system re-signs the server certificate passed to the client browser, then acts as a man-in-the-middle to decrypt the session. You can optionally replace the self-signed certificate key only and not the entire certificate, in which case users see a self-signed certificate key notice in the browser.

Controlling Traffic Based on Encrypted Session Characteristics

The system can control encrypted traffic based on the cipher suite or server certificate used to negotiate the session. You can configure one of several different reusable objects and reference the object in a decryption rule condition to match traffic. The following table describes the different types of reusable objects you can configure:

If you configure...	You can control encrypted traffic based on whether...
A cipher suite list containing one or more cipher suites	The cipher suite used to negotiate the encrypted session matches a cipher suite in the cipher suite list

If you configure...	You can control encrypted traffic based on whether...
A trusted CA object by uploading a CA certificate your organization trusts	The trusted CA trusts the server certificate used to encrypt the session, whether: • The CA issued the certificate directly • The CA issued a certificate to an intermediate CA that issued the server certificate
An external certificate object by uploading a server certificate	The server certificate used to encrypt the session matches the uploaded server certificate
A distinguished name object containing a certificate subject or issuer distinguished name	The subject or issuer common name, country, organization, or organizational unit on the certificate used to encrypt the session matches the configured distinguished name

Related Topics

Cipher Suite List, on page 1378
Distinguished Name, on page 1382
PKI, on page 1404

Decryption Rule Order Evaluation

This information applies only to rule-based decryption policies and rules.

When you create a decryption rule in a decryption policy, you specify its position using the **Insert** list in the rule editor. Decryption rules in an a decryption policy are numbered, starting at 1. The system matches traffic to decryption rules in top-down order by ascending rule number.

In most cases, the system handles network traffic according to the *first* decryption rule where *all* the rule's conditions match the traffic. Except in the case of Monitor rules (which log traffic but do not affect traffic flow), the system does *not* continue to evaluate traffic against additional, lower-priority rules after that traffic matches a rule. Conditions can be simple or complex; you can control traffic by security zone, network or geographical location, VLAN, port, application, requested URL category and reputation, user, certificate, certificate distinguished name, certificate status, cipher suite, or encryption protocol version.

Each rule also has an *action*, which determines whether you monitor, block, or inspect matching encrypted or decrypted traffic with access control. Note that the system does *not* further inspect encrypted traffic it blocks. It does subject encrypted and undecryptable traffic to access control. However, access control rule conditions require unencrypted traffic, so encrypted traffic matches fewer rules.

Rules that use *specific* conditions (such as network and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your rules. For more information about the OSI model, see this Wikipedia article.

Tip Proper decryption rule order reduces the resources required to process network traffic, and prevents rule preemption. Although the rules you create are unique to every organization and deployment, there are a few general guidelines to follow when ordering rules that can optimize performance while still addressing your needs.

In addition to ordering rules by number, you can group rules by category. By default the system provides three categories: Administrator, Standard, and Root. You can add custom categories, but you cannot delete the system-provided categories or change their order.

Related Topics

Best Practices for Access Control Rules, on page 1738
Default Handling Options for Undecryptable Traffic, on page 2296
Decryption Policy Rule Order

Decryption Rule Conditions

A decryption rule's conditions identify the type of encrypted traffic the rule handles. Conditions can be simple or complex, and you can specify more than one condition type per rule. Only if traffic meets all the conditions in a rule does the rule apply to the traffic.

If you do not configure a particular condition for a rule, the system does not match traffic based on that criterion. For example, a rule with a certificate condition but no version condition evaluates traffic based on the server certificate used to negotiate the session, regardless of the session SSL or TLS version.

Every decryption rule has an associated action that determines the following for matching encrypted traffic:

- Handling: Most importantly, the rule action governs whether the system will monitor, trust, block, or decrypt encrypted traffic that matches the rule's conditions
- Logging: The rule action determines when and how you can log details about matching encrypted traffic.

Your TLS/SSL inspection configuration handles, inspects, and logs decrypted traffic:

- The decryption policy's undecryptable actions handle traffic that the system cannot decrypt.
- The policy's default action handles traffic that does not meet the condition of any non-Monitor decryption rule.

You can log a connection event when the system blocks or trusts an encrypted session. You can also force the system to log connections that it decrypts for further evaluation by access control rules, regardless of how the system later handles or inspects the traffic. Connection logs for encrypted sessions contain details about the encryption, such as the certificate used to encrypt that session. You can log only end-of-connection events, however:

- For blocked connections (Block, Block with reset), the system immediately ends the sessions and generates an event
- For Do Not Decrypt connections, the system generates an event when the session ends

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Caution Adding the first or removing the last active authentication rule when TLS/SSL decryption is disabled (that is, when the access control policy does not include a decryption policy) restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

Note that an active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive or VPN identity cannot be established** selected.

Related Topics

Security Zone Rule Conditions, on page 1905
Network Rule Conditions, on page 939
VLAN Tags Rule Conditions, on page 1787
User Rule Conditions, on page 939
Application Rule Conditions, on page 939
Port Rule Conditions, on page 940
Category Rule Conditions, on page 2321
Server Certificate-Based Decryption Rule Conditions, on page 2321

Security Zone Rule Conditions

Security zones segment your network to help you manage, classify, and decrypt traffic flow by grouping interfaces across multiple devices.

Security zones control or decrypt traffic by its source and destination security zones. If you add both source and destination zones to a zone condition, matching traffic must originate from an interface in one of the source zones and leave through an interface in one of the destination zones.

Just as all interfaces in a zone must be of the same type (all inline, passive, switched, or routed), all zones used in a zone condition must be of the same type. Because devices deployed passively do not transmit traffic, you cannot use a zone with passive interfaces as a destination zone.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Tip Constraining rules by zone is one of the best ways to improve system performance. If a rule does not apply to traffic through any of device's interfaces, that rule does not affect that device's performance.

Security Zone Conditions and Multitenancy

In a multidomain deployment, a zone created in an ancestor domain can contain interfaces that reside on devices in different domains. When you configure a zone condition in an descendant domain, your configurations apply to only the interfaces you can see.

Network Rule Conditions

Networks control or decrypt traffic by its source and destination IP address, using inner headers. Tunnel rules, which use outer headers, have tunnel endpoint conditions instead of network conditions.

You can use predefined objects to build network conditions, or manually specify individual IP addresses or address blocks.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Note You *cannot* use FDQN network objects in identity rules.

VLAN Tags Rule Conditions

Note VLAN tags in access rules only apply to inline sets. Access rules with VLAN tags do not match traffic on firewall interfaces.

VLAN rule conditions control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic. The system uses the innermost VLAN tag to filter VLAN traffic, with the exception of the prefilter policy, which uses the outermost VLAN tag in its rules.

Note the following Q-in-Q support:

- Firewall Threat Defense on Firepower 4100/9300—Does not support Q-in-Q (supports only one VLAN tag).
- Firewall Threat Defense on all other models:
  - Inline sets and passive interfaces—Supports Q-in-Q, up to 2 VLAN tags.
  - Firewall interfaces—Does not support Q-in-Q (supports only one VLAN tag).

You can use predefined objects to build VLAN conditions, or manually enter any VLAN tag from 1 to 4094. Use a hyphen to specify a range of VLAN tags.

In a cluster, if you encounter problems with VLAN matching, edit the access control policy advanced options, Transport/Network Preprocessor Settings, and select the **Ignore the VLAN header when tracking connections** option.

User Rule Conditions

Matches traffic based the user who initiates the connection, or the group to which the user belongs. For example, you could configure a Block rule to prohibit anyone in the Finance group from accessing a network resource.

You can configure user rule conditions for users in Microsoft Active Directory realms only.

In addition to configuring users and groups for configured realms, you can set policies for the following Special Identities users:

- Failed Authentication: User that failed authentication with the captive portal.
- Guest: Users configured as guest users in the captive portal.
- No Authentication Required: Users that match an identity **No Authentication Required** rule action.
- Unknown: Users that cannot be identified; for example, users that are not downloaded by a configured realm.

For access control rules only, you must first associate an identity policy with the access control policy as discussed in Associating Other Policies with Access Control, on page 1762.

Client Threat Rule Conditions

The threat confidence levels of the client processes are determined by the Encrypted Visibility Engine (EVE), and represent the likelihood that the incoming traffic is malicious. The client threat confidence levels range from very low to very high, and the levels can be specified as rule conditions for decryption policy rules. When the client threat confidence levels and server URL category reputation are configured, they act as a powerful combination in decryption policies. We highly recommend that you specify a URL Category Reputation rule condition for rules that specify a Client Threat rule condition. When Client Threat and URL Category Reputation conditions are both specified for a decryption policy rule, the reputations of both the client and the server are utilized to make more accurate decryption policy decisions. We recommend that you do not specify a Client Threat rule condition alone without specifying a URL Category Reputation rule condition because this rule configuration only acts on information about the client and lacks information about the server.

Prerequisites for configuring client threat rule conditions

- The EVE must be enabled in the corresponding access control policy.
- The management center and threat defense devices must be Version 7.7.
- Snort 3 must be enabled on the device.
- Valid IPS and URL Filtering licenses must be installed on the managed device.

We recommend that you bypass decryption for very low-risk clients connecting to trusted servers based on the client threat determined by EVE and URL Category Reputation. To do so, create a new decryption policy rule, or edit an existing rule with the following configurations:

1. Choose the decryption policy rule action as **Do not decrypt**.
2. In the **Client Threat** tab, choose the client threat level as **Very Low**.
3. In the **Category** tab, choose the category as **All** and the reputation as **Trusted**.

Application Rule Conditions

When the system analyzes IP traffic, it can identify and classify the commonly used applications on your network. This discovery-based *application awareness* is the basis for *application control*—the ability to control application traffic.

System-provided *application filters* help you perform application control by organizing applications according to basic characteristics: type, risk, business relevance, category, and tags. You can create reuseable user-defined filters based on combinations of the system-provided filters, or on custom combinations of applications.

At least one detector must be enabled for each application rule condition in the policy. If no detector is enabled for an application, the system automatically enables all system-provided detectors for the application; if none exist, the system enables the most recently modified user-defined detector for the application. For more information about application detectors, see Application Detector Fundamentals, on page 2624.

You can use both application filters and individually specified applications to ensure complete coverage. However, understand the following note before you order your access control rules.

Benefits of Application Filters

Application filters help you quickly configure application control. For example, you can easily use system-provided filters to create an access control rule that identifies and blocks all high risk, low business relevance applications. If a user attempts to use one of those applications, the system blocks the session.

Using application filters simplifies policy creation and administration. It assures you that the system controls application traffic as expected. Because Cisco frequently updates and adds application detectors via system and vulnerability database (VDB) updates, you can ensure that the system uses up-to-date detectors to monitor application traffic. You can also create your own detectors and assign characteristics to the applications they detect, automatically adding them to existing filters.

Application Characteristics

The system characterizes each application that it detects using the criteria described in the following table. Use these characteristics as application filters.

Table 126: Application Characteristics

Characteristic	Description	Example
Type	Application protocols represent communications between hosts. Clients represent software running on a host. Web applications represent the content or requested URL for HTTP traffic.	HTTP and SSH are application protocols. Web browsers and email clients are clients. MPEG video and Facebook are web applications.
Risk	The likelihood that the application is being used for purposes that might be against your organization's security policy.	Peer-to-peer applications tend to have a very high risk.
Business Relevance	The likelihood that the application is being used within the context of your organization's business operations, as opposed to recreationally.	Gaming applications tend to have a very low business relevance.
Category	A general classification for the application that describes its most essential function. Each application belongs to at least one category.	Facebook is in the social networking category.
Tag	Additional information about the application. Applications can have any number of tags, including none.	Video streaming web applications often are tagged `high bandwidth` and `displays ads`.

Related Topics

Best Practices for Configuring Application Control, on page 1735

Port Rule Conditions

Port conditions allow you to control traffic by its source and destination ports.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Best Practices for Port-Based Rules

Specifying ports is the traditional way to target applications. However, applications can be configured to use unique ports to bypass access control blocks. Thus, whenever possible, use application filtering criteria rather than port criteria to target traffic.

Application filtering is also recommended for applications, like Firewall Threat Defense, that open separate channels dynamically for control vs. data flow. Using port-based access control rules can prevent these kinds of applications from performing correctly, and could result in blocking desirable connections.

Using Source and Destination Port Constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as source port conditions in a single access control rule.

Category Rule Conditions

You can optionally choose to include categories in your decryption policies. These categories, also referred to as *URL filtering*, are updated by the Cisco Talos intelligence group. Updates are based on machine learning and human analysis according to content that is retrievable from the website destination and sometimes from its hosting and registration information. Categorization is *not* based on the declared company vertical, intent, or security.

For more information, see URL Filtering Overview, on page 1857.

If you are using category rule conditions in decryption policies in a rule with the **Do Not Decrypt** rule action, see Decryption Rule Do Not Decrypt Action, on page 2335.

Server Certificate-Based Decryption Rule Conditions

decryption rules can handle and decrypt encrypted traffic based on server certificate characteristics. You can configure decryption rules based on the following server certificate attributes:

- Distinguished name conditions allow you to handle and inspect encrypted traffic based on the CA that issued a server certificate, or the certificate holder. Based on the issuer distinguished name, you can handle traffic based on the CA that issued a site's server certificate.

- Certificate conditions in decryption rules allow you to handle and inspect encrypted traffic based on the server certificate used to encrypt that traffic. You can configure a condition with one or more certificates; traffic matches the rule if the certificate matches any of the condition's certificates.
- Certificate status conditions in decryption rules allow you to handle and inspect encrypted traffic based on the status of the server certificate used to encrypt the traffic, including whether a certificate is valid, revoked, expired, not yet valid, self-signed, signed by a trusted CA, whether the Certificate Revocation List (CRL) is valid; whether the Server Name Indication (SNI) in the certificate matches the server in the request.
- Cipher suite conditions in decryption rules allow you to handle and inspect encrypted traffic based on the cipher suite used to negotiate the encrypted session.
- Session conditions in decryption rules allow you to inspect encrypted traffic based on the SSL or TLS version used to encrypt the traffic.

To detect multiple cipher suites in a rule, the certificate issuer, or the certificate holder, you can create reusable cipher suite list and distinguished name objects and add them to your rule. To detect the server certificate and certain certificate statuses, you must create external certificate and external CA objects for the rule.

Related Topics

Certificate Decryption Rule Conditions, on page 2322
Certificate Status Decryption Rule Conditions, on page 2329
Trusting External Certificate Authorities, on page 2327
Matching Traffic on Certificate Status
Cipher Suite Decryption Rule Conditions, on page 2331
Encryption Protocol Version Decryption Rule Conditions, on page 2334

Certificate Decryption Rule Conditions

When you build a certificate-based decryption rule condition, you can upload a server certificate; you save the certificate as an external certificate *object*, which is reusable and associates a name with a server certificate. Alternately, you can configure certificate conditions with existing external certificate objects and object groups.

You can search the **Available Certificates** field in the rule condition based for external certificate objects and object groups based on the following certificate distinguished name characteristics:

- Subject or issuer common name (CN), or if the URL is contained in the certificate's Subject Alternative Name (SAN)

  The URL the user enters in the browser matches the Common Name (CN)
- Subject or issuer organization (O)
- Subject or issuer organizational unit (OU)

You can choose to match against multiple certificates in a single certificate rule condition; if the certificate used to encrypt the traffic matches any of the uploaded certificates, the encrypted traffic matches the rule.

You can add a maximum of 50 external certificate objects and external certificate object groups to the **Selected Certificates** in a single certificate condition.

Note the following:

- You cannot configure a certificate condition if you also select the **Decrypt - Known Key** action. Because that action requires you to select a server certificate to decrypt traffic, the implication is that the certificate already matches the traffic.
- If you configure a certificate condition with an external certificate object, any cipher suites you add to a cipher suite condition, or internal CA objects you associate with the **Decrypt - Resign** action, must match the external certificate's signature algorithm type. For example, if your rule's certificate condition references an EC-based server certificate, any cipher suites you add, or CA certificates you associate with the **Decrypt - Resign** action, must also be EC-based. If you mismatch signature algorithm types in this case, the policy editor displays a warning next to the rule.
- The first time the system detects an encrypted session to a new server, certificate data is not available for ClientHello processing, which can result in an undecrypted first session. After the initial session, the managed device caches data from the server Certificate message. For subsequent connections from the same client, the system can match the ClientHello message conclusively to rules with certificate conditions and process the message to maximize decryption potential.

Distinguished Name (DN) Rule Conditions

This topic discusses how to use distinguished name conditions in a decryption rule. If you're not sure, you can find a certificate's Subject Alternative Name (SAN) and Common Name using a web browser, then you can add those values to a decryption rule as distinguished name conditions.

For detailed information about SANs, see RFC 528, section 4.2.1.6.

The following sections discuss:

- DN rule matching example
- How the system uses the SNIs and SANs
- How to find a certificate's Common Name and Subject Alternative Names
- How to add a DN rule condition

DN rule matching example

Following is an example of DN rule conditions in a Do Not Decrypt rule. Suppose you want to make sure to *not* decrypt traffic going to `amp.cisco.com` or to YouTube. You could set up your DN conditions as follows:

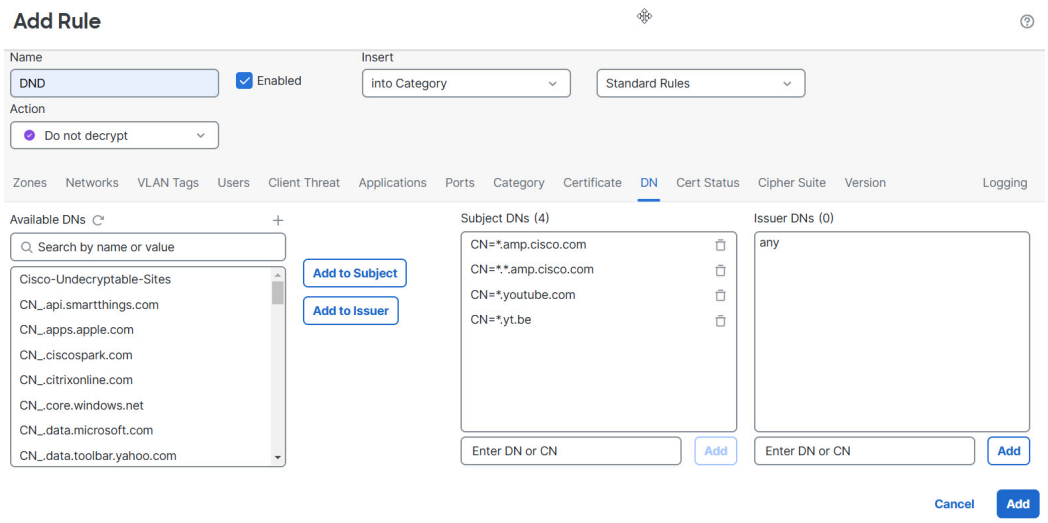

The preceding DN rule conditions would match the following URLs and therefore, the traffic would be undecrypted an earlier rule prevented it:

- www.amp.cisco.com
- auth.amp.cisco.com
- auth.us.amp.cisco.com
- www.youtube.com
- kids.youtube.com
- www.yt.be

The preceding DN rule conditions would *not* match any of the following URLs and therefore, the traffic would not match the Do Not Decrypt rule but might match any other decryption rules in the same decryption policy.

- amp.cisco.com
- youtube.com
- yt.be

To match any of the preceding host names, add more CNs to the rule (for example, adding `CN=yt.be` would match that URL.)

How the system uses the SNI and SANs

The host name portion of the URL in the client request is the Server Name Indication (SNI). The client specifies which hostname they want to connect to (for example, `auth.amp.cisco.com`) using the SNI extension in the TLS handshake. The server then selects the corresponding private key and certificate chain that are required to establish the connection while hosting all certificates on a single IP address.

If there's a match between the SNI and the CN or a SAN in the certificate, we use the SNI when comparing against the DNs listed in the rule. If there is no SNI or if it doesn't match the certificate, we use the certificate's CN when comparing against the DNs listed in the rule.

How to find a certificate's Common Name and subject alternative names

To find any certificate's Common Name, use the following steps. You can even use these steps to find the common name and SANs for a self-signed certificate.

These steps are for Firefox but other browsers are similar. The following procedure uses `amp.cisco.com` as an example.

1. Browse to `amp.cisco.com` in Firefox.
2. In the browser's location bar, to the left of the URL, click .
3. Click **Connection secure** > **More Information**.

   (For a non-secure or self-signed certificate, click **Connection not secure** > **More Information**.)
4. On the Page Info dialog box, click **View Certificate**.

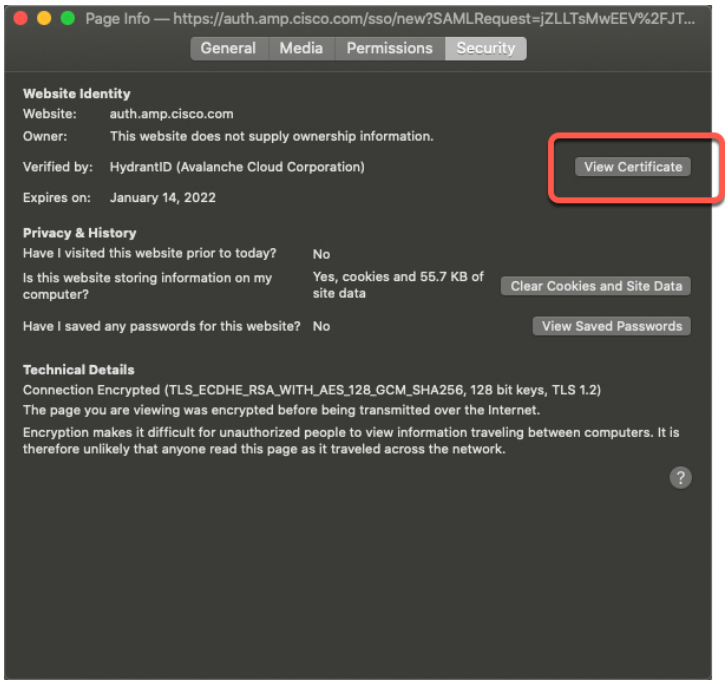

5. The next page shows certificate details.

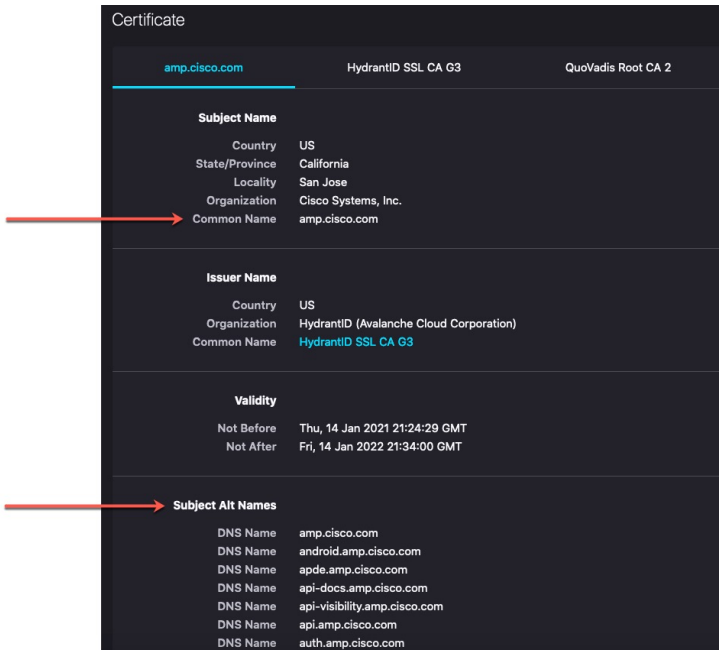

Note the following:

- `CN=auth.amp.cisco.com`, if used as a DN rule condition, would match *only* that host name (that is, SNI). The SNI `amp.cisco.com` would *not* match.
- To match as many domain name fields as possible, use wildcards.

  For example, to match `auth.amp.cisco.com`, use `CN=*.amp.cisco.com`. To match `auth.us.amp.cisco.com`, use `CN=*.*.amp.cisco.com`.

  A DN like `CN=*.example.com` matches `www.example.com` but *not* `example.com`. To match both SNIs, use two DNs in the rule condition.
- Don't go overboard with wildcards though. For example, a DN object like `CN=*.google.com` matches a very large number of SANs. Instead of `CN=*.google.com`, use a DN object like `CN=*.youtube.com` as the DN object so it matches names like `www.youtube.com`.

  You can also use variations of the SNI that match SANs like `CN=*.youtube.com`, `CN=youtu.be`, `CN=*.yt.be`, and so on.
- A self-signed certificate should work the same way. You can confirm it's a self-signed certificate by the fact the issuer DN is the same as the subject DN.

How to add a DN rule condition

After you know the CN you want to match, edit the decryption rule in one of the following ways:

- Use an existing DN.

  Click the name of a DN and then click either **Add to Subject** or **Add to Issuer**. (**Add to Subject** is much more common.) To view the value of a DN object, hover the mouse pointer over it.)

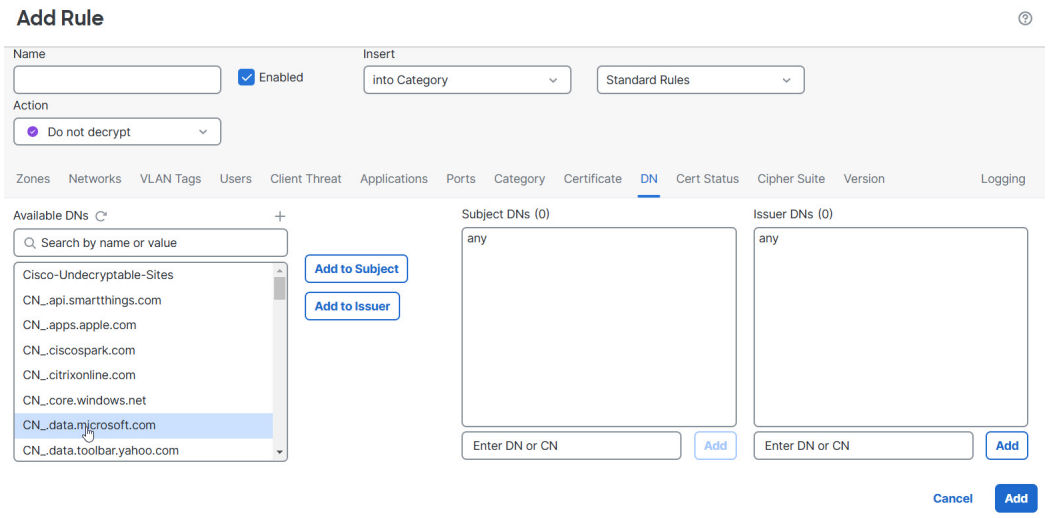

- Create a new DN object.

  Click **Add** (+) to the right of Available DNs. The DN object must consist of a name and a value.

- Add the DN directly.

  Enter the DN in the field at the bottom of the **Subject DNs** field or the **Issuer DNs** field. (**Subject DNs** is more common.) After you enter the DN, click **Add**.

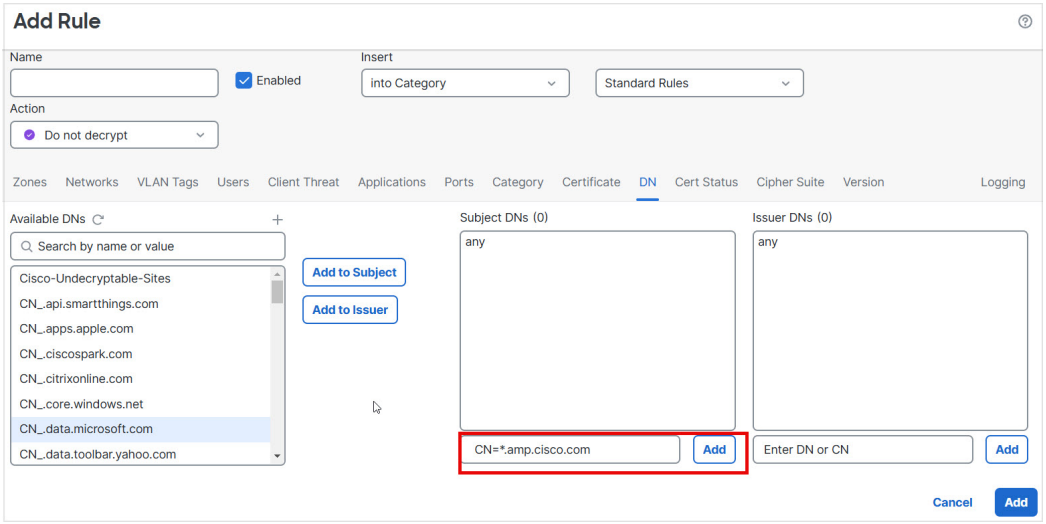

Related Topics

Distinguished Name, on page 1382

Trusting External Certificate Authorities

You can trust CAs by adding root and intermediate CA certificates to your decryption policy, then use these trusted CAs to verify server certificates used to encrypt traffic.

If a trusted CA certificate contains an uploaded certificate revocation list (CRL), you can also verify whether a trusted CA revoked the encryption certificate.

Tip Upload all certificates in a root CA's chain of trust to the list of trusted CA certificates, including the root CA certificate and all intermediate CA certificates. Otherwise, it is more difficult to detect trusted certificates issued by intermediate CAs. Also, if you configure certificate status conditions to trust traffic based on the root issuer CA, all traffic within a trusted CA's chain of trust can be allowed without decryption, rather than unnecessarily decrypting it.

For more information, see Trusted CA Object, on page 1409.

Note When you create a decryption policy, the policy's **Trusted CA Certificate** tab page is populated with several trusted CA certificates, including the **Cisco-Trusted-Authorities** group, which is added to the **Select Trusted CAs** list.

Procedure

Step 1 Log in to the Firewall Management Center if you haven't already done so.

Step 2 Click **Policies** > **Access Control heading** > **Decryption**.

Step 3 Click **Edit** (✎) next to the decryption policy to edit.

Step 4 Click **Add Rule** to add a new decryption rule or click **Edit** (✎) to edit an existing rule.

Step 5 Click the **Certificates** tab.

Step 6 Find the trusted CAs you want to add from the **Available Certificates**, as follows:

- To add a trusted CA object on the fly, which you can then add to the condition, click **Add** (+) above the **Available Certificates** list.
- To search for trusted CA objects and groups to add, click the **Search by name or value** prompt above the **Available Certificates** list, then enter either the name of the object, or a value in the object. The list updates as you type to display matching objects.

Step 7 To select an object, click it. To select all objects, right-click and then **Select All**.

Step 8 Click **Add to Rule**.

Tip
You can also drag and drop selected objects.

Step 9 Add or continue editing the rule.

What to do next

- Add a certificate status decryption rule condition to your SSL rule. See Matching Traffic on Certificate Status for more information.

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Certificate Status Decryption Rule Conditions

For each certificate status decryption rule you configure, you can match traffic against the presence or absence of a given status. You can select several statuses in one rule condition; if the certificate matches any of the selected statuses, the rule matches the traffic.

You can choose to match against the presence or absence of multiple certificate statuses in a single certificate status rule condition; the certificate needs to match only one of the criteria to match the rule.

You should consider, when setting this parameter, whether you're configuring a decrypt rule or a block rule. Typically, you should click **Yes** for a block rule and **No** for a decrypt rule. Examples:

- If you're configuring a **Decrypt - Resign** rule, the default behavior is to decrypt traffic with an expired certificate. To change that behavior, click **No** for **Expired** so traffic with an expired certificate is not decrypted and resigned.
- If you're configuring a **Block** rule, the default behavior is to allow traffic with an expired certificate. To change that behavior click **Yes** for **Expired** so traffic with an expired certificate is blocked.

The following table describes how the system evaluates encrypted traffic based on the encrypting server certificate's status.

Table 127: Certificate Status Rule Condition Criteria

Status Check	Status Set to Yes	Status Set to No
Revoked	The policy trusts the CA that issued the server certificate, and the CA certificate uploaded to the policy contains a CRL that revokes the server certificate.	The policy trusts the CA that issued the certificate, and the CA certificate uploa policy does not contain a CRL that rev certificate.
Self-signed	The detected server certificate contains the same subject and issuer distinguished name.	The detected server certificate contains subject and issuer distinguished names
Valid	All of the following are true: • The policy trusts the CA that issued the certificate. • The signature is valid. • The issuer is valid. • None of the policy's trusted CAs revoked the certificate. • The current date is between the certificate Valid From and Valid To date.	At least one of the following is true: • The policy does not trust the CA t certificate. • The signature is invalid. • The issuer is invalid. • A trusted CA in the policy revoke certificate. • The current date is before the cert From date. • The current date is after the certifi date.
Invalid signature	The certificate's signature cannot be properly validated against the certificate's content.	The certificate's signature is properly val the certificate's content.

Status Check	Status Set to Yes	Status Set to No
Invalid issuer	The issuer CA certificate is not stored in the policy's list of trusted CA certificates.	The issuer CA certificate is stored in the p of trusted CA certificates.
Expired	The current date is after the certificate Valid To date.	The current date is before or on the certific To date.
Not yet valid	The current date is before the certificate Valid From date.	The current date is after or on the certificat From date.
Invalid certificate	The certificate is not valid. At least one of the following is true: • Invalid or inconsistent certificate extension; that is, a certificate extension had an invalid value (for example, an incorrect encoding) or some value inconsistent with other extensions. • The certificate cannot be used for the specified purpose. • The Basic Constraints path length parameter has been exceeded. For more information, see RFC 5280, section 4.2.1.9. • The certificate's value for Not Before or Not After is invalid. These dates can be encoded as UTCTime or GeneralizedTime For more information, see RFC 5280 section 4.1.2.5. • The format of the name constraint is not recognized; for example, an email address format of a form not mentioned in RFC 5280, section 4.2.1.10. This could be caused by an improper extension or some new feature not currently supported. An unsupported name constraint type was encountered. OpenSSL currently supports only directory name, DNS name, email, and URI types. • The root certificate authority is not trusted for the specified purpose. • The root certificate authority rejects the specified purpose.	The certificate is valid. All of the followin • Valid certificate extension. • The certificate can be used for the spe purpose. • Valid Basic Constraints path length. • Valid values for Not Before and Not A • Valid name constraint. • The root certificate is trusted for the s purpose. • The root certificate accepts the specifie

Status Check	Status Set to Yes	Status Set to No
Invalid CRL	The Certificate Revocation List (CRL) digital signature is not valid. At least one of the following is true: • The value of the CRL's Next Update or Last Update field is invalid. • The CRL is not yet valid. • The CRL has expired. • An error occurred when attempting to verify the CRL path. This error occurs only if extended CRL checking is enabled. • CRL could not be found. • The only CRLs that could be found did not match the scope of the certificate.	The CRL is valid. All of the following • Next Update and Last Update fiel • The CRL's date is valid. • The path is valid. • The CRL was found. • The CRL matches the certificate's
Server mismatch	The server name does not match the server's Server Name Indication (SNI) name, which could indicate an attempt to spoof the server name.	The server name matches the SNI name to which the client is requesting access

Note that even though a certificate might match more than one status, the rule causes an action to be taken on the traffic only once.

Checking whether a CA issued or revoked a certificate requires uploading root and intermediate CA certificates and associated CRLs as objects. You then add these trusted CA objects to an decryption policy's list of trusted CA certificates.

Cipher Suite Decryption Rule Conditions

The system provides predefined cipher suites you can add to a cipher suite rule condition for Block or Block with Reset rule actions. You can also add cipher suite list objects containing multiple cipher suites.

Important Cipher suite rule conditions can be used only to *block* traffic, they cannot be used to *decrypt* traffic.

Note You cannot add new cipher suites. You can neither modify nor delete predefined cipher suites.

You can add a maximum of 50 cipher suites and cipher suite lists to the **Selected Cipher Suites** in a single cipher suite condition. The system supports adding the following cipher suites to a cipher suite condition:

- SSL_RSA_FIPS_WITH_3DES_EDE_CBC_SHA
- SSL_RSA_FIPS_WITH_DES_CBC_SHA
- TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA

- TLS_DHE_RSA_WITH_AES_128_CBC_SHA
- TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
- TLS_DHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_DHE_RSA_WITH_AES_256_CBC_SHA
- TLS_DHE_RSA_WITH_AES_256_CBC_SHA256
- TLS_DHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA
- TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA256
- TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA
- TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA256
- TLS_DHE_RSA_WITH_DES_CBC_SHA
- TLS_DH_Anon_WITH_AES_128_GCM_SHA256
- TLS_DH_Anon_WITH_AES_256_GCM_SHA384
- TLS_DH_Anon_WITH_CAMELLIA_128_CBC_SHA
- TLS_DH_anon_WITH_CAMELLIA_128_CBC_SHA256
- TLS_DH_Anon_WITH_CAMELLIA_256_CBC_SHA
- TLS_DH_anon_WITH_CAMELLIA_256_CBC_SHA256
- TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384
- TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_ECDSA_WITH_NULL_SHA
- TLS_ECDHE_ECDSA_WITH_RC4_128_SHA
- TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384

- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_RSA_WITH_NULL_SHA
- TLS_ECDHE_RSA_WITH_RC4_128_SHA
- TLS_RSA_WITH_3DES_EDE_CBC_SHA
- TLS_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_AES_128_CBC_SHA256
- TLS_RSA_WITH_AES_128_GCM_SHA256
- TLS_RSA_WITH_AES_256_CBC_SHA
- TLS_RSA_WITH_AES_256_CBC_SHA256
- TLS_RSA_WITH_AES_256_GCM_SHA384
- TLS_RSA_WITH_CAMELLIA_128_CBC_SHA
- TLS_RSA_WITH_CAMELLIA_128_CBC_SHA256
- TLS_RSA_WITH_CAMELLIA_256_CBC_SHA
- TLS_RSA_WITH_CAMELLIA_256_CBC_SHA256
- TLS_RSA_WITH_DES_CBC_SHA
- TLS_RSA_WITH_NULL_MD5
- TLS_RSA_WITH_NULL_SHA
- TLS_RSA_WITH_RC4_128_MD5
- TLS_RSA_WITH_RC4_128_SHA

Note the following:

- If you add cipher suites not supported for your deployment, you cannot deploy your configuration. For example, passive deployments do not support decrypting traffic with the any of the ephemeral Diffie-Hellman (DHE) or ephemeral elliptic curve Diffie-Hellman (ECDHE) cipher suites. Creating a rule with these cipher suites prevents you from deploying your access control policy.
- You can add an anonymous cipher suite to the **Cipher Suite** condition in decryption policy to use the rule, you must also configure your in an order that prevents ClientHello processing. For more information, see Decryption Policy Rule Order.
- When specifying a cipher suite as a rule condition, consider that the rule matches on the negotiated cipher suite in the ServerHello message, rather than on the full list of cipher suites specified in the ClientHello message. During ClientHello processing, the managed device strips unsupported cipher suites from the ClientHello message. However, if this results in all specified cipher suites being stripped, the system retains the original list. If the system retains unsupported cipher suites, subsequent evaluation results in an undecrypted session.

Encryption Protocol Version Decryption Rule Conditions

You can choose to match against traffic encrypted with SSL version 3.0, or TLS version 1.0, 1.1, or 1.2. By default, all protocol versions are selected when you create a rule; if you select multiple versions, encrypted traffic that matches any of the selected versions matches the rule. You must select at least one protocol version when saving the rule condition.

You can use SSL 3.0 in a Do Not Decrypt, Block, or Block with Reset rule action.

You *cannot* select SSL v2.0 in a version rule condition; the system does not support decrypting traffic encrypted with SSL version 2.0. You can configure an undecryptable action to allow or block this traffic without further inspection. For more information, see Set Default Handling for Undecryptable Traffic, on page 2297.

Important Protocol version rule conditions can be used *only* to *block* traffic, they cannot be used to *decrypt* traffic.

For example, to block all SSL v3.0, TLS v1.0, TLS v1.1, and TLS v1.2 traffic, set the options as follows:

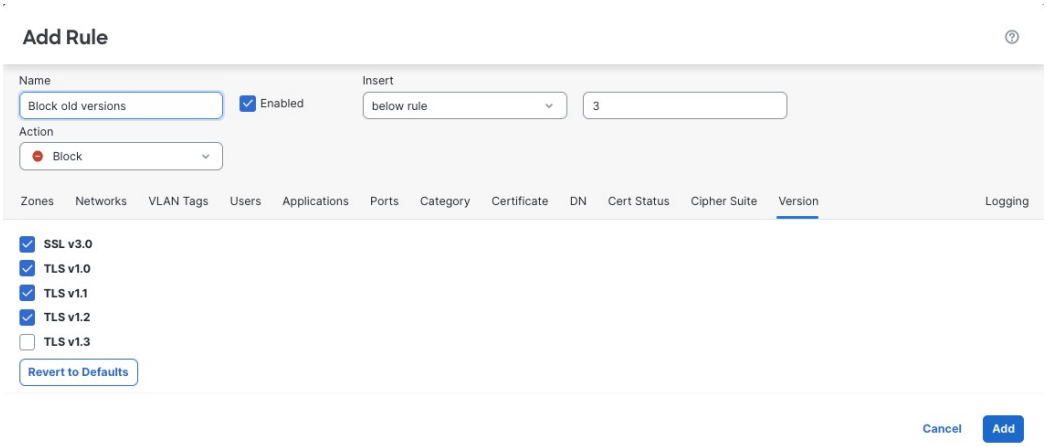

Decryption Rule Actions

The following sections discuss the actions available with Decryption rules.

Decryption Rule Monitor Action

The **Monitor** action is not designed to permit or deny traffic. Rather, its primary purpose is to force connection logging, regardless of how matching traffic is eventually handled. The ClientHello message is not modified if traffic matches a **Monitor** rule condition.

Traffic is then matched against additional rules, if present, to determine whether to trust, block, or decrypt it. The first non-Monitor rule matched determines traffic flow and any further inspection. If there are no additional matching rules, the system uses the default action.

Because the primary purpose of Monitor rules is to track network traffic, the system automatically logs end-of connection events for monitored traffic to the Secure Firewall Management Center database, regardless of the logging configuration of the rule or default action that later handles the connection.

Decryption Rule Do Not Decrypt Action

The **Do Not Decrypt** action passes encrypted traffic for evaluation by the access control policy's rules and default action. Because some access control rule conditions require unencrypted traffic, this traffic might match fewer rules. The system cannot perform deep inspection on encrypted traffic, such as intrusion or file inspection.

Typical reasons for a **Do Not Decrypt** rule action include:

- When decrypting TLS/SSL traffic is prohibited by law.
- Sites you know you can trust.
- Sites you can disrupt by inspecting traffic (such as Windows Update).
- To view the values of TLS/SSL fields using connection events. (You do not need to decrypt traffic to view connection event fields.) For more information, see *Requirements for Populating Connection Event Fields* in the Cisco Secure Firewall Management Center Administration Guide.

For more information, see Default Handling Options for Undecryptable Traffic, on page 2296

Limitations of categories in Do Not Decrypt rules

You can optionally choose to include categories in your decryption policies. These categories, also referred to as *URL filtering*, are updated by the Cisco Talos intelligence group. Updates are based on machine learning and human analysis according to content that is retrievable from the website destination and sometimes from its hosting and registration information. Categorization is *not* based on the declared company vertical, intent, or security. While we strive to continuously update and improve URL filtering categories, it is not an exact science. Some websites are not categorized at all and it's possible some websites might be improperly categorized.

Avoid overusing categories in do not decrypt rules to avoid decrypting traffic without a reason; for example, the Health and Medicine category includes the WebMD website, which does not threaten patient privacy.

Following is a sample decryption policy that can prevent decryption for websites in the Health and Medicine categories but allows decryption for WebMD and everything else. General information about decryption rules can be found in Guidelines for Using TLS/SSL Decryption, on page 2304.

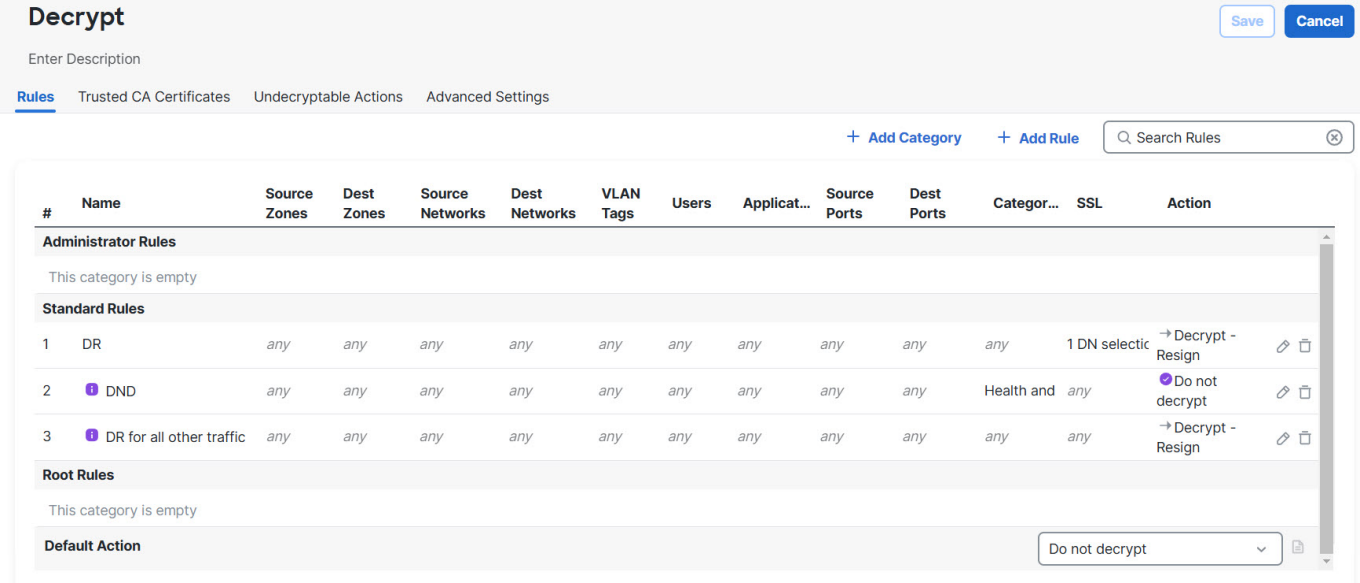

Note Don't confuse URL filtering with application detection, which relies on reading some of the packet from a website to determine more specifically what it is (for example, Facebook Message or Salesforce). For more information, see Best Practices for Configuring Application Control, on page 1735.

Decryption Rule Blocking Actions

The system provides the following decryption rule actions for traffic you do not want to pass through the system:

- **Block** to terminate the connection, resulting in an error in the client browser.

  The error message does not indicate the site was blocked due to policy. Instead, errors might indicate that there are no common encryption algorithms. It is not obvious from this message that you blocked the connection on purpose.

- **Block with reset** to terminate and reset the connection, resulting in an error in the client browser.

  The error indicates the connection was reset but does not indicate why.

Tip You cannot use the **Block** or **Block with reset** action in a passive or inline (tap mode) deployment because the device does not directly inspect the traffic. If you create a rule with the **Block** or **Block with reset** action that contains passive or inline (tap mode) interfaces within a security zone condition, the policy editor displays a warning (⚠) next to the rule.

Decryption Rule Decrypt Actions

The **Decrypt - Replace Cert**, **Decrypt - Known Key**, and **Decrypt - Resign** actions decrypt encrypted traffic. The system inspects decrypted traffic with access control. Access control rules handle decrypted and unencrypted traffic identically — you can inspect it for discovery data as well as detect and block intrusions, prohibited files, and malware. The system reencrypts allowed traffic before passing it to its destination.

We recommend you use a certificate from a trusted Certificate Authority (CA) to decrypt traffic. This prevents **Invalid Issuer** from being displayed in the SSL Certificate Status column in connection events.

For more information about adding trusted objects, see Trusted Certificate Authority Objects, on page 1409.

Related topic: TLS 1.3 Decryption Best Practices, on page 2300.

Related Topics

TLS 1.3 Decryption Best Practices, on page 2300

Select Internal Certificate Objects

Before you begin

To create a decryption rule with a **Decrypt - Known Key** action, you must add internal certificates for the servers to decrypt traffic. For more information, see Internal Certificate Authority Objects, on page 1405.

Procedure

Step 1 Click **Edit** (✎) in the field next to the decryption rule **Action**.

Step 2 In the Select Internal Certificate Objects dialog box, do the following:
Available Certificates list: Click the name of a certificate and click **Add to Rule** to use it in this rule. You can add one or more certificates to the rule.

Selected Certificates list: Displays internal certificates used in the rule. To remove a certificate, click **Delete** (🗑).

Step 3 To save these changes, click **OK**.

Monitor TLS/SSL Hardware Acceleration

The following topics discuss how to monitor the status of TLS/SSL

Informational Counters

If a system under load is working well, you should see large counts for the following counters. Because there are 2 sides to the tracker process per connection, you can see these counters increase by 2 per connection. The PRIV_KEY_RECV and SECU_PARAM_RECV counters are the most important, and are highlighted. The CONTEXT_CREATED and CONTEXT_DESTROYED counters relate to the allocation of cryptographic chip memory.

```
> show counters
Protocol      Counter                          Value   Context
SSLENC        CONTEXT_CREATED                 258225   Summary
SSLENC        CONTEXT_DESTROYED               258225   Summary
TLS_TRK       OPEN_SERVER_SESSION             258225   Summary
TLS_TRK       OPEN_CLIENT_SESSION             258225   Summary
TLS_TRK       UPSTREAM_CLOSE                  516450   Summary
TLS_TRK       DOWNSTREAM_CLOSE                516450   Summary
TLS_TRK       FREE_SESSION                    516450   Summary
TLS_TRK       CACHE_FREE                      516450   Summary
TLS_TRK       PRIV_KEY_RECV                   258225   Summary
TLS_TRK       NO_KEY_ENABLE                   258225   Summary
TLS_TRK       SECU_PARAM_RECV                 516446   Summary
TLS_TRK       DECRYPTED_ALERT                 258222   Summary
TLS_TRK       DECRYPTED_APPLICATION         33568976   Summary
TLS_TRK       ALERT_RX_CNT                    258222   Summary
TLS_TRK       ALERT_RX_WARNING_ALERT          258222   Summary
TLS_TRK       ALERT_RX_CLOSE_NOTIFY           258222   Summary
TCP_PRX       OPEN_SESSION                    516450   Summary
TCP_PRX       FREE_SESSION                    516450   Summary
TCP_PRX       UPSTREAM_CLOSE                  516450   Summary
TCP_PRX       DOWNSTREAM_CLOSE                516450   Summary
TCP_PRX       FREE_CONN                       258222   Summary
TCP_PRX       SERVER_CLEAN_UP                 258222   Summary
TCP_PRX       CLIENT_CLEAN_UP                 258222   Summary
```

Alert Counters

We implemented the following counters according to the TLS 1.2 specification. FATAL or BAD alerts could indicate issues; however, ALERT_RX_CLOSE_NOTIFY is normal.

For details, see RFC 5246 section 7.2.

```
TLS_TRK       ALERT_RX_CNT                       311   Summary
TLS_TRK       ALERT_TX_CNT                         2   Summary
TLS_TRK       ALERT_TX_IN_HANDSHAKE_CNT            2   Summary
TLS_TRK       ALERT_RX_IN_HANDSHAKE_CNT            2   Summary
TLS_TRK       ALERT_RX_WARNING_ALERT             308   Summary
TLS_TRK       ALERT_RX_FATAL_ALERT                 3   Summary
TLS_TRK       ALERT_TX_FATAL_ALERT                 2   Summary
TLS_TRK       ALERT_RX_CLOSE_NOTIFY              308   Summary
TLS_TRK       ALERT_RX_BAD_RECORD_MAC              2   Summary
TLS_TRK       ALERT_TX_BAD_RECORD_MAC              2   Summary
TLS_TRK       ALERT_RX_BAD_CERTIFICATE             1   Summary
```

Error Counters

These counters indicate system errors. These counts should be low on a healthy system. The BY_PASS counters indicate packets that have been passed directly to or from the inspection engine (Snort) process (which runs in software) without decryption. The following example lists some of the bad counters.

Counters with a value of 0 are not displayed. To view a complete list of counters, use the command **show counters description | include TLS_TRK**

```
> show counters
Protocol      Counter                          Value   Context
TCP_PRX       BYPASS_NOT_ENOUGH_MEM             2134   Summary
TLS_TRK       CLOSED_WITH_INBOUND_PACKET           2   Summary
TLS_TRK       ENC_FAIL                            82   Summary
```

```
TLS_TRK      DEC_FAIL                          211   Summary
TLS_TRK      DEC_CKE_FAIL                    43194   Summary
TLS_TRK      ENC_CB_FAIL                      4335   Summary
TLS_TRK      DEC_CB_FAIL                       909   Summary
TLS_TRK      DEC_CKE_CB_FAIL                   818   Summary
TLS_TRK      RECORD_PARSE_ERR                  123   Summary
TLS_TRK      IN_ERROR                        44948   Summary
TLS_TRK      ERROR_UPSTREAM_RECORD           43194   Summary
TLS_TRK      INVALID_CONTENT_TYPE              123   Summary
TLS_TRK      DOWNSTREAM_REC_CHK_ERROR          123   Summary
TLS_TRK      DECRYPT_FAIL                    43194   Summary
TLS_TRK      UPSTREAM_BY_PASS                  127   Summary
TLS_TRK      DOWNSTREAM_BY_PASS                127   Summary
```

Fatal Counters

The fatal counters indicate serious errors. These counters should be at or near 0 on a healthy system. The following example lists the fatal counters.

```
> show counters
Protocol     Counter                         Value   Context
CRYPTO       RING_FULL                           1   Summary
CRYPTO       ACCELERATOR_CORE_TIMEOUT            1   Summary
CRYPTO       ACCELERATOR_RESET                   1   Summary
CRYPTO       RSA_PRIVATE_DECRYPT_FAILED          1   Summary
```

The RING_FULL counter is not a fatal counter, but indicates how often the system overloaded the cryptographic chip. The ACCELERATOR_RESET counter is the number of times the TLS crypto acceleration process failed unexpectedly, which also causes the failure of pending operations, which are the numbers you see in ACCELERATOR_CORE_TIMEOUT and RSA_PRIVATE_DECRYPT_FAILED.

If you have persistent problems, disable TLS crypto acceleration (or **config hwCrypto disable**) and work with Cisco TAC to resolve the issues.

Note You can do additional troubleshooting using the **show snort tls-offload** and **debug snort tls-offload** commands. Use the **clear snort tls-offload** command to reset the counters displayed in the **show snort tls-offload** command to zero.

Troubleshoot Decryption Rules

The following topics discuss how to troubleshoot decryption rules.

About TLS/SSL Oversubscription

TLS/SSL oversubscription is a state where a managed device is overloaded with TLS/SSL traffic. Any managed device can experience TLS/SSL oversubscription but only managed devices that support TLS crypto acceleration provide a configurable way to handle it.

When a managed device with TLS crypto acceleration enabled is oversubscribed, any packet received by the managed device is acted on according to the setting for **Handshake Errors** in the decryption policy's **Undecryptable Actions**:

- Inherit default action
- Do not decrypt
- Block
- Block with reset

If the setting for **Handshake Errors** in the decryption policy's **Undecryptable Actions** is **Do Not decrypt** and the associated access control policy is configured to inspect the traffic, inspection occurs; decryption does *not* occur.

Troubleshoot TLS/SSL Oversubscription

If your managed device has TLS crypto acceleration enabled, you can view connection events to determine whether or not the devices are experiencing SSL oversubscription. You must add at least the **SSL Flow Flags** event to the table view of connection events.

Before you begin

- Configure a decryption policy with a setting for **Handshake Errors** on **Undecryptable Actions** page.

  For more information, see Set Default Handling for Undecryptable Traffic, on page 2297.

- Enable logging for your SSL rules as discussed in the section on logging decryptable connections in decryption rules in the Secure Firewall Management Center and Threat Defense Management Network Administration guide.

Procedure

Step 1 If you haven't done so already, log in to the Firewall Management Center.

Step 2 Click **Analysis** > **Connections**.

Step 3 Click **Table View of Connection Events**.

Step 4 Click **x** on any column in the connection events table to add additional columns for at least **SSL Flow Flags** and **SSL Flow Messages**.

The following example shows adding the **SSL Actual Action**, **SSL Flow Error**, **SSL Flow Flags**, **SSL Flow Messages**, **SSL Policy**, and **SSL Rule** columns to the table of connection events. (Look in the Disabled Columns section of the dialog box.)

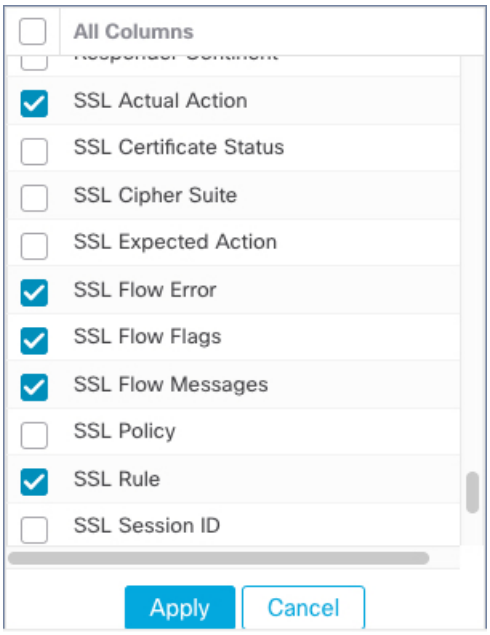

The columns are added in the order discussed in *Connection and Security Intelligence Event Fields* in the Cisco Secure Firewall Management Center Administration Guide .

Step 5 Click **Apply**.
TLS/SSL oversubscription is indicated by the values of `ERROR_EVENT_TRIGGERED` and `OVER_SUBSCRIBED` in the **SSL Flow Flags** column.

Step 6 If TLS/SSL oversubscription is occurring, log in to the managed device and enter any of the following commands:

Command	Result
show counters	If the value of **TCP_PRX BYPASS_NOT_ENOUGH_MEM** is large, consider upgrading your device to one with a larger capacity for SSL traffic or use **Do Not Decrypt** rules for lower priority encrypted traffic.
show snort tls-offload	If the value of **BYPASS_NOT_ENOUGH_MEM** is large, consider upgrading your device to one with a larger capacity for SSL traffic or use **Do Not Decrypt** rules for lower priority encrypted traffic.

About TLS Heartbeat

Some applications use the *TLS heartbeat* extension to the Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) protocols defined by RFC6520. TLS heartbeat provides a way to confirm the connection is still alive—either the client or server sends a specified number of bytes of data and requests the other party echo the response. If this is successful, encrypted data is sent.

When a managed device with TLS crypto acceleration enabled encounters a packet that uses the TLS heartbeat extension, the managed device takes the action specified by the setting for **Decryption Errors** in the decryption policy's **Undecryptable Actions**:

- Block
- Block with reset

Related Topics

Troubleshoot TLS Heartbeat, on page 2342

Troubleshoot TLS Heartbeat

If your managed device has TLS crypto acceleration enabled, you can view connection events to determine whether or not the devices are seeing traffic with the TLS heartbeat extension. You must add at least the **SSL Flow Messages** event to the table view of connection events.

Before you begin

SSL heartbeat is indicated by the value of `HEARTBEAT` in the **SSL Flow Messages** column in the table view of connection events. To determine if applications in your network use SSL heartbeat, first perform the following tasks:

- Configure an decryption policy with a setting for **Decryption Errors** on **Undecryptable Actions** page.

  For more information, see Set Default Handling for Undecryptable Traffic, on page 2297.

- Enable logging for your SSL rules as discussed in Secure Firewall Management Center and Threat Defense Management Network Administration.

Procedure

Step 1 If you haven't done so already, log in to the Firewall Management Center.

Step 2 Click **Analysis** > **Connections**.

Step 3 Click **Table View of Connection Events**.

Step 4 Click **x** on any column in the connection events table to add additional columns for at least **SSL Flow Flags** and **SSL Flow Messages**.

The following example shows adding the **SSL Actual Action**, **SSL Flow Error**, **SSL Flow Flags**, **SSL Flow Messages**, **SSL Policy**, and **SSL Rule** columns to the table of connection events.

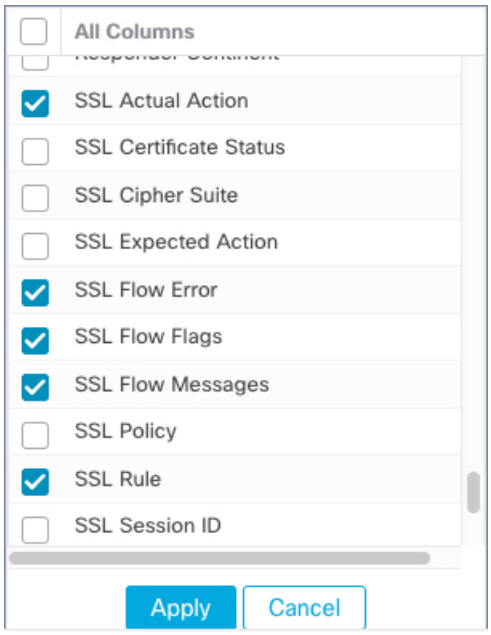

The columns are added in the order discussed in *Connection and Security Intelligence Event Fields* in the Cisco Secure Firewall Management Center Administration Guide .

Step 5 Click **Apply**.
TLS heartbeat is indicated by the value of `HEARTBEAT` in the **SSL Flow Messages** column.

Step 6 If applications in your network use SSL heartbeat, see Decryption Rule Guidelines and Limitations, on page 2304.

About TLS/SSL Pinning

Some applications use a technique referred to as *TLS/SSL pinning* or *certificate pinning*, which embeds the fingerprint of the original server certificate in the application itself. As a result, if you configured a decryption rule with a **Decrypt - Resign** action, when the application receives a resigned certificate from a managed device, validation fails and the connection is aborted.

To confirm that TLS/SSL pinning is occurring, attempt to log in to a mobile application like Facebook. If a network connection error is displayed, log in using a web browser. (For example, you *cannot* log in to a Facebook mobile application but *can* log in to Facebook using Safari or Chrome.) You can use Firepower Management Center connection events as further proof of TLS/SSL pinning

Note TLS/SSL pinning is not limited to mobile applications.

If applications in your network use SSL pinning, see TLS/SSL Certificate Pinning Guidelines, on page 2310.

Related Topics

Troubleshoot TLS/SSL Pinning, on page 2344

Troubleshoot TLS/SSL Pinning

You can view connection events to determine whether or not the devices are experiencing SSL pinning. You must add at least the **SSL Flow Flags** and **SSL Flow Messages** columns to the table view of connection events.

Before you begin

- Enable logging for your decryption rules as discussed in the section on logging decryptable connections in decryption rules in the Secure Firewall Management Center and Threat Defense Management Network Administration guide.
- Log in to a mobile application like Facebook; if a network connection error displays, log in to Facebook using Chrome or Safari. If you *can* log in using a web browser but not the native application, SSL pinning is likely occurring.

Procedure

Step 1 If you haven't done so already, log in to the Firewall Management Center.

Step 2 Click **Analysis** > **Connections**.

Step 3 Click **Table View of Connection Events**.

Step 4 Click **x** on any column in the connection events table to add additional columns for at least **SSL Flow Flags** and **SSL Flow Messages**.

The following example shows adding the **SSL Actual Action**, **SSL Flow Error**, **SSL Flow Flags**, **SSL Flow Messages**, **SSL Policy**, and **SSL Rule** columns to the table of connection events.

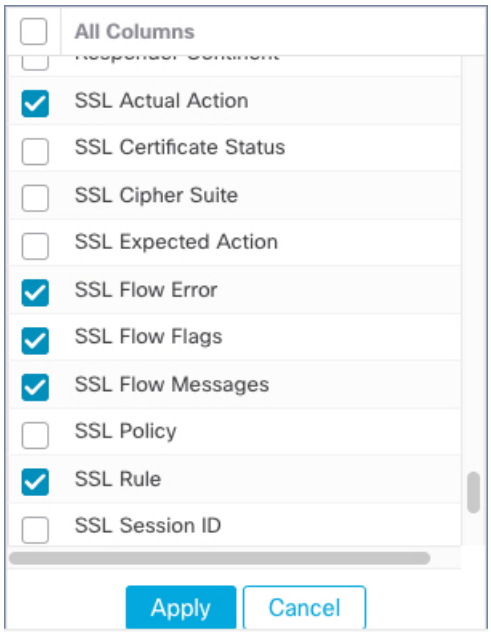

The columns are added in the order discussed in the section on connection and security intelligence event fields in the Secure Firewall Management Center and Threat Defense Management Network Administration guide.

Step 5 Click **Apply**.

Step 6 The following paragraphs discuss how you can identify SSL pinning behavior.

Step 7 If you determine that applications in your network use SSL pinning, see Decryption Rule Guidelines and Limitations, on page 2304.

What to do next

You can use TLS/SSL connection events to confirm TLS/SSL pinning is occurring by looking for any of the following:

- Applications that send an SSL ALERT Message as soon as the client receives the SERVER_HELLO, SERVER_CERTIFICATE, SERVER_HELLO_DONE message from the server, followed by a TCP Reset, exhibit the following symptoms. (The alert, `Unknown CA (48)`, can be viewed using a packet capture.)
  - The SSL Flow Flags column displays `ALERT_SEEN` but *not* `APP_DATA_C2S` or `APP_DATA_S2C`.
  - If your managed device has SSL hardware acceleration enabled, the SSL Flow Messages column typically displays: `CLIENT_ALERT`, `CLIENT_HELLO`, `SERVER_HELLO`, `SERVER_CERTIFICATE`, `SERVER_KEY_EXCHANGE`, `SERVER_HELLO_DONE`.
  - If your managed device doesn't support SSL hardware acceleration or if the feature is disabled, the SSL Flow Messages column typically displays: `CLIENT_HELLO`, `SERVER_HELLO`, `SERVER_CERTIFICATE`, `SERVER_KEY_EXCHANGE`, `SERVER_HELLO_DONE`.
  - `Success` is displayed in the SSL Flow Error column.

- Applications that send no alerts but instead send TCP Reset after the SSL handshake is finished exhibit the following symptoms:
  - The SSL Flow Flags column does *not* display `ALERT_SEEN`, `APP_DATA_C2S`, or `APP_DATA_S2C`.
  - If your managed device has SSL hardware acceleration enabled, the SSL Flow Messages column typically displays: `CLIENT_HELLO`, `SERVER_HELLO`, `SERVER_CERTIFICATE`, `SERVER_KEY_EXCHANGE`, `SERVER_HELLO_DONE`, `CLIENT_KEY_EXCHANGE`, `CLIENT_CHANGE_CIPHER_SPEC`, `CLIENT_FINISHED`, `SERVER_CHANGE_CIPHER_SPEC`, `SERVER_FINISHED`.
  - If your managed device doesn't support SSL hardware acceleration or if the feature is disabled, the SSL Flow Messages column typically displays: `CLIENT_HELLO`, `SERVER_HELLO`, `SERVER_CERTIFICATE`, `SERVER_KEY_EXCHANGE`, `SERVER_HELLO_DONE`, `CLIENT_KEY_EXCHANGE`, `CLIENT_CHANGE_CIPHER_SPEC`, `CLIENT_FINISHED`,`SERVER_CHANGE_CIPHER_SPEC`, `SERVER_FINISHED`.
  - `Success` is displayed in the SSL Flow Error column.

Related Topics

Troubleshoot Unknown or Bad Certificates or Certificate Authorities, on page 2346

Troubleshoot Unknown or Bad Certificates or Certificate Authorities

You can view connection events to determine whether or not the devices are experiencing unknown certificate authorities, bad certificates, or unknown certificates. This procedure can also be used if a TLS/SSL certificate has been pinned. You must add at least the **SSL Flow Flags** and **SSL Flow Messages** columns to the table view of connection events.

Before you begin

- Set up a decryption rule.
- Enable logging for your decryption rules as discussed in the section on logging decryptable connections in decryption rules in the Secure Firewall Management Center and Threat Defense Management Network Administration guide.

Procedure

Step 1 If you haven't done so already, log in to the Firewall Management Center.

Step 2 Click **Analysis** > **Connections**.

Step 3 Click **Table View of Connection Events**.

Step 4 Click **x** on any column in the connection events table to add additional columns for at least **SSL Flow Flags** and **SSL Flow Messages**.

The following example shows adding the **SSL Actual Action**, **SSL Flow Error**, **SSL Flow Flags**, **SSL Flow Messages**, **SSL Policy**, and **SSL Rule** columns to the table of connection events.

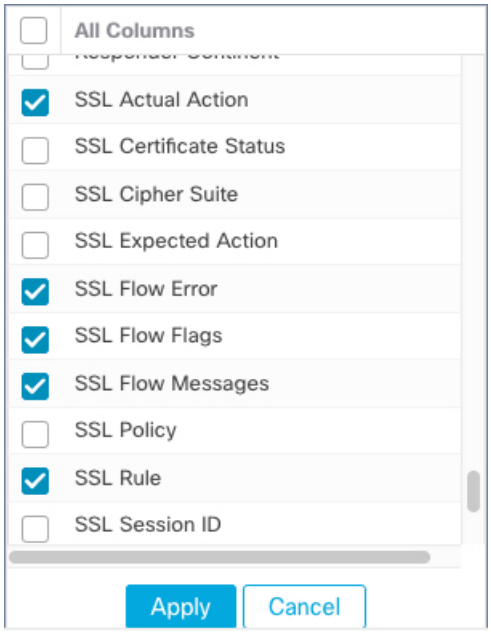

The columns are added in the order discussed in the section on connection and security intelligence event fields in the Secure Firewall Management Center and Threat Defense Management Network Administration guide.

Step 5 Click **Apply**.

Step 6 The following table discusses how you can determine if a certificate or certificate authority is bad or missing.

SSL flow flag	Meaning
CLIENT_ALERT_SEEN_UNKNOWN_CA	Indicates a valid certificate chain or partial chain was received by an SSL client application, but the certificate was not accepted because the CA certificate could not be located or could not be matched with a known, trusted CA. This message always indicates an unrecoverable error.
CLIENT_ALERT_SEEN_BAD_CERTIFICATE	A certificate was corrupt, contained signatures that did not verify correctly, or had other problems.
CLIENT_ALERT_SEEN_CERTIFICATE_UNKNOWN	Some other (unspecified) issue arose in processing the certificate, rendering it unacceptable.

Verify TLS/SSL Cipher Suites

Before you begin

This topic discusses actions you must take if you see the following error when saving a decryption rule that has cipher suite conditions:

```
Traffic cannot match this rule; none of your selected cipher suites contain a signature
algorithm that matches the resigning CA's signature algorithm
```

The error indicates that one or more of the cipher suites you chose for the decryption rule condition are incompatible with the certificate used in the decryption rule. To resolve the issue, you must have access to the certificate you're using.

Note The tasks in this topic assume knowledge of how TLS/SSL encryption works.

Procedure

Step 1 When you attempt to save a decryption rule rule with either **Decrypt - Resign**, **Decrypt - Replace Cert**, or **Decrypt - Known Key** with specified cipher suites, the following error is displayed:

Example:

```
Traffic cannot match this rule; none of your selected cipher suites contain a
signature algorithm that the resigning CA's signature algorithm
```

Step 2 Locate the certificate you're using to decrypt traffic and, if necessary, copy the certificate to a system that can run openssl commands.

Step 3 Run the following command to display the signature algorithm used by the certificate:

openssl x509 -in *CertificateName* **-text -noout**

The first few lines of output are displayed similar to the following:

```
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 4105 (0x1009)
    Signature Algorithm: ecdsa-with-SHA256
```

Step 4 The **Signature algorithm** tells you the following:

- The cryptographic function used (in the preceding example, **ECDSA** means Elliptic Curve Digital Signature Algorithm).
- The hash function used to create a digest of the encrypted message (in the preceding example, **SHA256**).

Step 5 Search a resource such as *OpenSSL at University of Utah* for cipher suites that match those values. The cipher suite must be in RFC format.
You can also search a variety of other sites, such as Server Side TLS at the Mozilla wiki or Appendix C of RFC 5246. Cipher Suites in TLS/SSL (Schannel SSP) in Microsoft documentation has a detailed explanation of cipher suites.

Step 6 If necessary, translate the OpenSSL name to an RFC name that the Firepower Management System uses.

See the RFC mapping list on the on the `https://testssl.sh` site.

Step 7 The previous example, **ecdsa-with-SHA256**, can be found in the Modern Compatibility List on the Mozilla wiki.

a) Choose only cipher suites that have **ECDSA** and **SHA-256** in the name. These cipher suites follow:

```
ECDHE-ECDSA-AES128-GCM-SHA256
ECDHE-ECDSA-AES128-SHA256
```

b) Find the corresponding RFC cipher suite on RFC mapping list. These cipher suites follow:

```
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
```

Step 8 Add the preceding cipher suites to your decryption rule.

Troubleshooting Using Crypto Archives

About Crypto Archives

Crypto issues are difficult to triage. Crypto archives help you to troubleshoot these issues. A crypto archive contains crypto session information about the crypto request, peer information, the component that sent the crypto request, and the timed-out crypto session information. Threat defense does not save keys and initialization vectors (IVs) for the session. For SSL and IPsec, you can also view the following information:

- For SSL: Session SSL version, source, destination IP addresses, and ports.
- For IPsec: IPsec security association information.

A ring can hold 2000 crypto command entries. Threat defense pushes the crypto command in one of the rings and pulls out the result after completing the crypto request. Crypto archive files now have the timed-out crypto request's ring and entry index. The ring and its entry index help in troubleshooting problematic crypto commands.

The crypto archive has two formats: a text file and a binary file. You can use the **debug menu ctm 103** command to decode the binary file. The crypto archive files are available at disk0:/crypto_archive.

For example:

```
FTD# debug menu ctm 103 crypto_eng0_arch_4.bin
[Nitrox V Archive Header v1.0 Info]
ASA Image Version: PIX (9.20) #0: Tue Mar 29 16:20:30 GMT 2022
...
SE SSL microcode: CNN5x-MC-SE-SSL-0011
AE microcode: CNN5x-MC-AE-MAIN-0002
Crypto Engine 0
Crash type: SE Ring Timeout
...
Core Soft Resets: 11
...
Timeout Ring (SE): 12
Timeout Entry: 642
SE TIMEOUT:
Core SE 6 Touts: 2
Core SE 8 Touts: 2
Core SE 12 Touts: 4
Core SE 32 Touts: 2
```

```
Core SE 37 Touts: 1
.....
[Timeout Session Info]
Active: TRUE
Sync: FALSE
Callback: TRUE
Saved Callback: FALSE
Commands in progress: 1
Engine : hardware
Device : n5 (Nitrox V)
Session : ssl
Priority: normal
NP VPN context handle : 0x00000000
Flag : 0
vcid : 0
Block size : 2050
async cb ring index: 0
tls offload rsa: FALSE
Session context:
SSL Version : dtls1.2
SSL Context Type : handshake
Encryption Mode : gcm
Auth Algorithm : null
Hash Algorithm : none
Key Size : 32
SSL V : dtls1.2
Source IP : 82.1.2.2
Source Port : 51915
Dest IP : 82.29.155.32
Dest Port : 443
```

In the above example, the highlighted information shows the timeout ring, the crash time (timeout entry), and SSL session information.

Supported Devices for Crypto Archives

The following devices with Nitrox V crypto accelerator support crypto archives:

- Cisco Firepower 3105, 3110, 3120, 3130, 3140
- Cisco Firepower 4112, 4115, 4125, 4145
- Cisco Firepower 9300 SM-40, SM-48 and SM-56
- Cisco Secure Firewall 4200

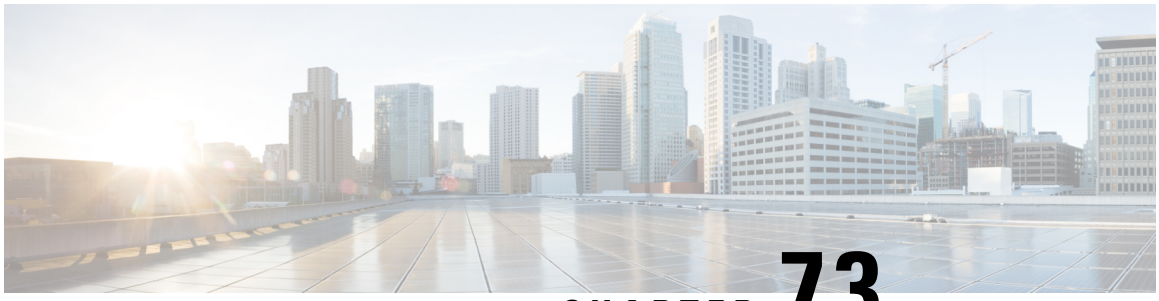

CHAPTER 73

Decryption Rules and Policy Example

This chapter builds on concepts discussed in this guide to provide a specific example of an SSL policy with decryption rules that follow our best practices and recommendations. You should be able to apply this example to your situation, adapting it to the needs of your organization.

In short:

- For trusted traffic (such as transferring a large compressed server backup), bypass inspection entirely, using prefiltering and flow offload.
- Put *first* any decryption rules that can be evaluated quickly, such as those that apply to specific IP addresses.
- Put *last* any decryption rules that require processing, **Decrypt - Resign**, and rules that block unsecure protocol versions and cipher suites.

- Decryption Rule Examples, on page 2351
- Run the Decryption Policy Wizard, on page 2351
- First Manual Do Not Decrypt Rule: Specific Traffic, on page 2356
- Next Manual Rule: Decrypt Specific Test Traffic, on page 2358
- Last Manual Decryption Rules: Block or Monitor Certificates and Protocol Versions, on page 2359
- Associate the Decryption Policy with an Access Control Policy and Advanced Settings, on page 2366
- Traffic to Prefilter, on page 2368
- Decryption Rule Settings, on page 2368

Decryption Rule Examples

This section provides an example of decryption rule that illustrate our best practices.

See one of the following sections for more information.

Run the Decryption Policy Wizard

This task discusses how to run the decryption policy wizard for outbound traffic protection. This policy has four rules:

1. **Do Not Decrypt** rule for distinguished names that are known to be undecryptable, most likely because they use TLS/SSL pinning.

2. **Do Not Decrypt** rule for URL categories that we classify as sensitive based on their content (for example, medical and financial).
3. **Do Not Decrypt** rule for applications that are known to be undecryptable, most likely because they use TLS/SSL pinning.
4. **Decrypt - Resign** rule that uses a certificate authority object named **IntCA** to decrypt the remainder of the traffic.

You can then edit the rules if you want and also manually add:

- **Decrypt - Resign** rules for traffic to monitor and determine whether the traffic should be blocked in the future.
- **Do Not Decrypt** rules for other types of traffic
- **Block** or **Block with Reset** rules for bad certificates and unsecure cipher suites.

If you enabled Change Management, you must create and assign a ticket before you can create a decryption policy. Before the decryption policy can be used, the ticket and all associated objects (like certificate authorities) must be approved. For more information, see Creating Change Management Tickets, on page 240 and Policies and Objects that Support Change Management, on page 235.

Procedure

Step 1 Log in to the Secure Firewall Management Center if you haven't already done so.

Step 2 Click **Policies** > **Access Control heading** > **Decryption**.

Step 3 Click **Create Decryption Policy**.

Step 4 Give the decryption policy a **Name** and optionally a **Description**.

Step 5 Click the **Outbound Protection** tab.

Step 6 From the **Internal CA** list, click the name of an internal certificate authority object or click **Create New** to upload or generate one.

The following figure shows an example.

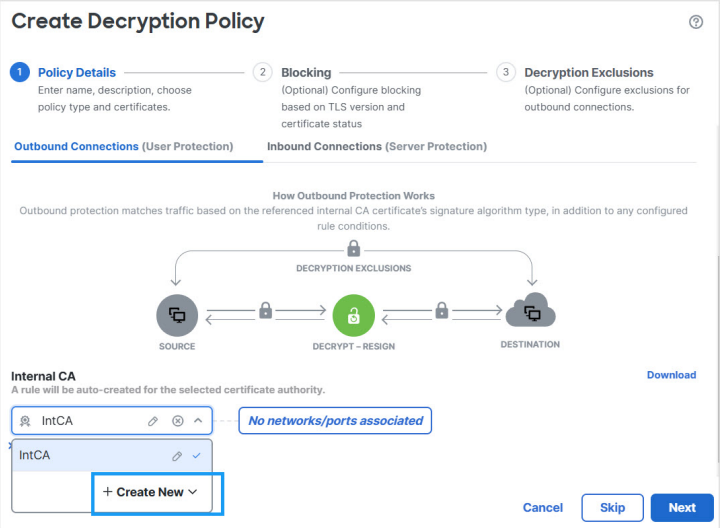

For more information about creating or uploading an internal certificate authority object, see:

- Upload an Internal CA for Outbound Protection, on page 2292
- Generate an Internal CA for Outbound Protection, on page 2291

Step 7 (Optional.) To restrict traffic to source and destination networks, click **Click to assign networks and ports**.

Step 8 Click **Next**.

Step 9 Complete the wizard as discussed in Decryption Policy Exclusions, on page 2287.

Decryption Policy Exclusions

This task discusses how to exclude certain types of traffic from decryption. We create **Do Not Decrypt** rules in your decryption policy for these although the rules are initially enabled only for an outbound decryption policy (that is, one that uses the **Decrypt - Resign** policy action).

Before you begin

You must upload an internal CA certificate for your managed device before you can create a decryption policy that protects outbound connections. You can do this in any of the following ways:

- Create an internal CA certificate object by going to **Objects** > **Object Management** > **PKI** > **Internal CAs** and referring to PKI, on page 1404.
- At the time you create this decryption policy.

Procedure

Step 1 Complete the tasks discussed in:

- Create a Create a Decryption Policy with Outbound Connection Protection, on page 2281

- For more information, see Create a Create a Decryption Policy with Inbound Connection Protection, on page 2285

Step 2 The exclusions page provides the following options. All options are *enabled* for an outbound protection policy (**Decrypt - Resign** rule action) and *disabled* for all other decryption policy actions.

Item	Description
Bypass decryption for sensitive URL categories	Check the box to not decrypt traffic from the indicated categories. Depending on the laws in your area, decryption certain traffic, such as finance or health-related, might be prohibited. Consult an authority in your area for more information. Click **Add** to add more categories. Click **Delete** (✕) to remove categories.
Bypass decryption for undecryptable distinguished names	Check the box to not decrypt traffic when re-signing the certificate is likely to cause the connection to fail. Typically, this behavior is associated with *certificate pinning*, which is discussed in TLS/SSL Certificate Pinning Guidelines, on page 2310. The list of undecryptable distinguished names is maintained by Cisco.
Bypass decryption for undecryptable applications	Check the box to not decrypt traffic when re-signing the certificate is likely to cause the connection to fail. Typically, this behavior is associated with *certificate pinning*, which is discussed in TLS/SSL Certificate Pinning Guidelines, on page 2310. Undecryptable applications are updated automatically in the Vulnerability Database (VDB). You can find a list of all applications on the Secure Firewall Application Detectors page; the **undecryptable** tag identifies applications Cisco determines are undecryptable. The list of undecryptable applications is maintained by Cisco.
Bypass decryption for very low-risk connections	Check this check box to not decrypt traffic for very low-risk clients connecting to trusted servers, based on the threat confidence levels of clients which is determined by the Encrypted Visibility Engine (EVE) and the URL Category Reputation. An **Auto-Rule-Low-Risk-Connections** Do Not Decrypt rule is created and enabled in the new decryption policy. To use the **Bypass decryption for very low-risk connections** option, you must enable the EVE in the corresponding access control policy, and the managed devices must be running Version 7.7 or later with a valid IPS license.

The following figure shows the default options.

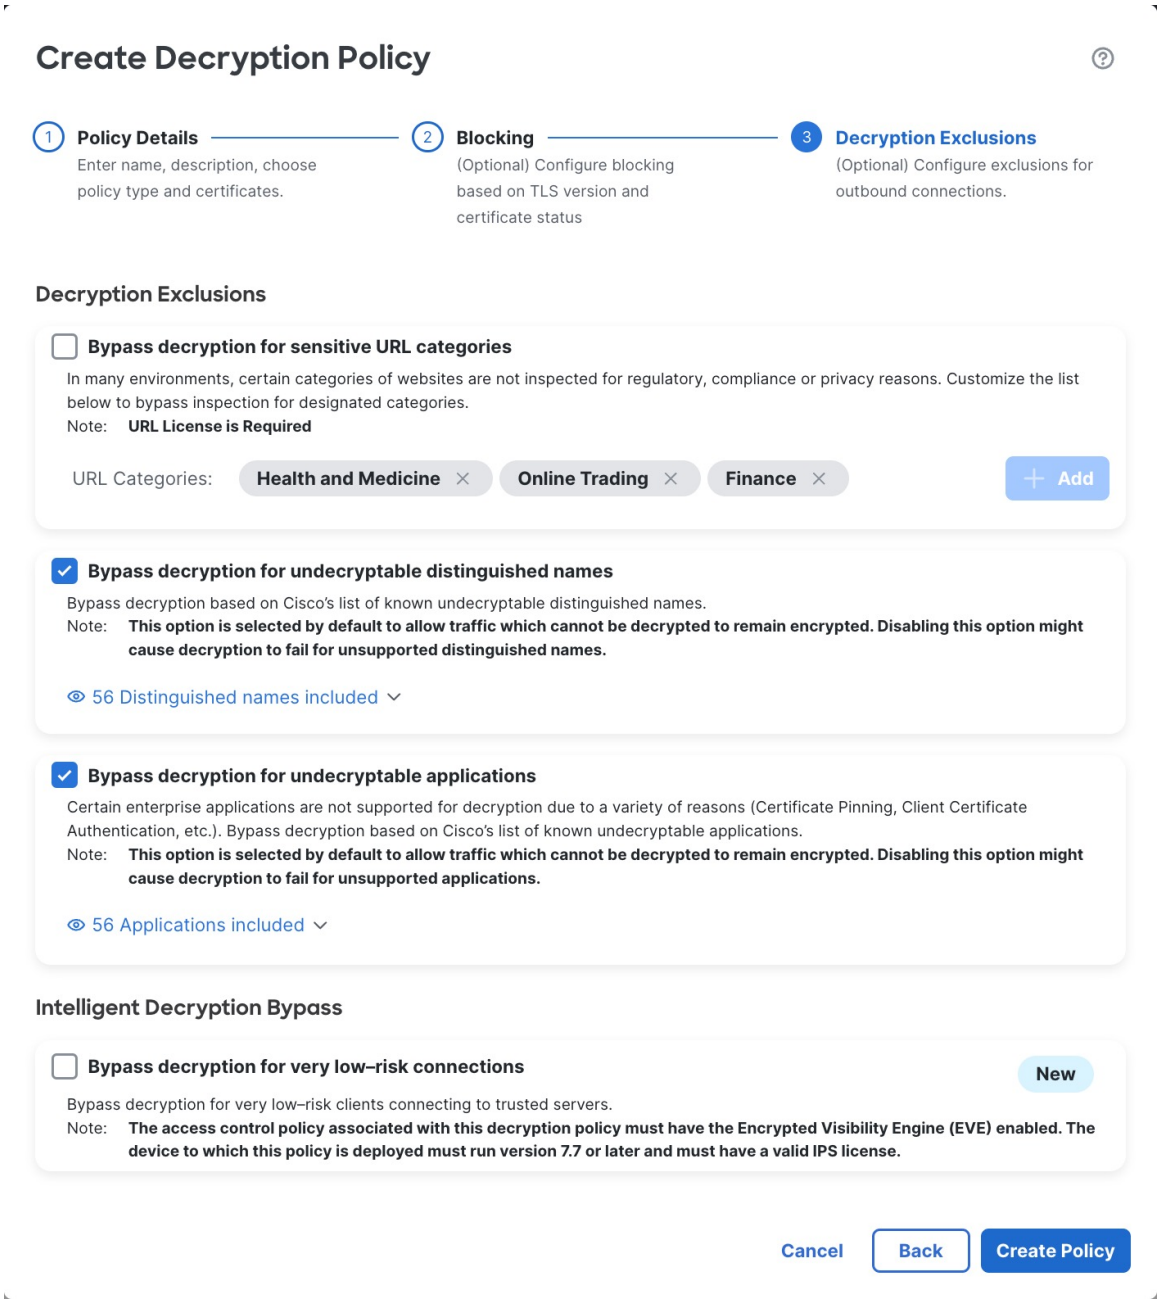

Step 3 Click **Create Policy**.

The following figure shows a sample outbound protection policy.

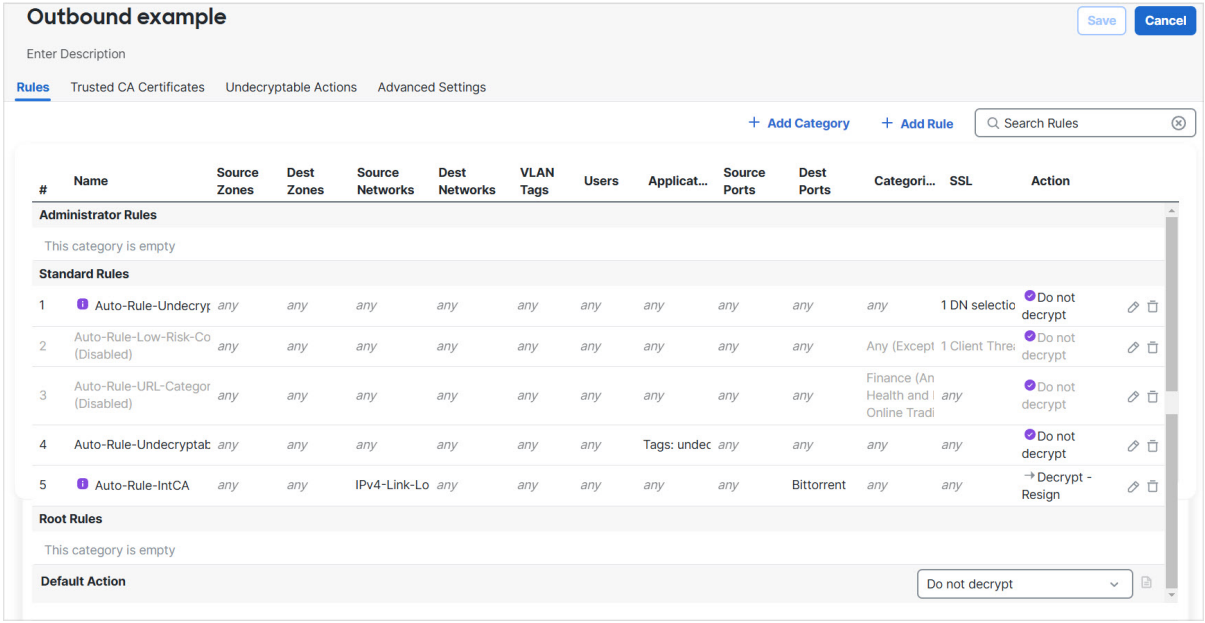

In the preceding example, the **Do Not Decrypt** rules corresponding to your choices for rule exclusions are automatically added before the **Decrypt - Resign** rule. The rule for sensitive URL categories is disabled because, by default, that exclusion is disabled. Had you selected the **Bypass decryption for sensitive URL categories** check box, the rule would have been enabled.

Step 4 Click **Create Policy**.

What to do next

- Add rule conditions: Decryption Rule Conditions, on page 2316
- Add a default policy action: Decryption Policy Default Actions, on page 2295
- Configure logging options for the default action as described in *Logging Connections with a Policy Default Action* in the Cisco Secure Firewall Management Center Administration Guide .
- Set advanced policy properties: Decryption Policy Advanced Options, on page 2298.
- Associate the decryption policy with an access control policy as described in Associating Other Policies with Access Control, on page 1762.
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

First Manual Do Not Decrypt Rule: Specific Traffic

The first decryption rule in the example does not decrypt traffic that goes to an internal network (defined as **internal**). **Do Not Decrypt** rule actions are matched during ClientHello so they are processed very fast.

After you run the decryption policy wizard, edit the policy to add the following rule. Drag it to the top of the list of rules so it's evaluated first.

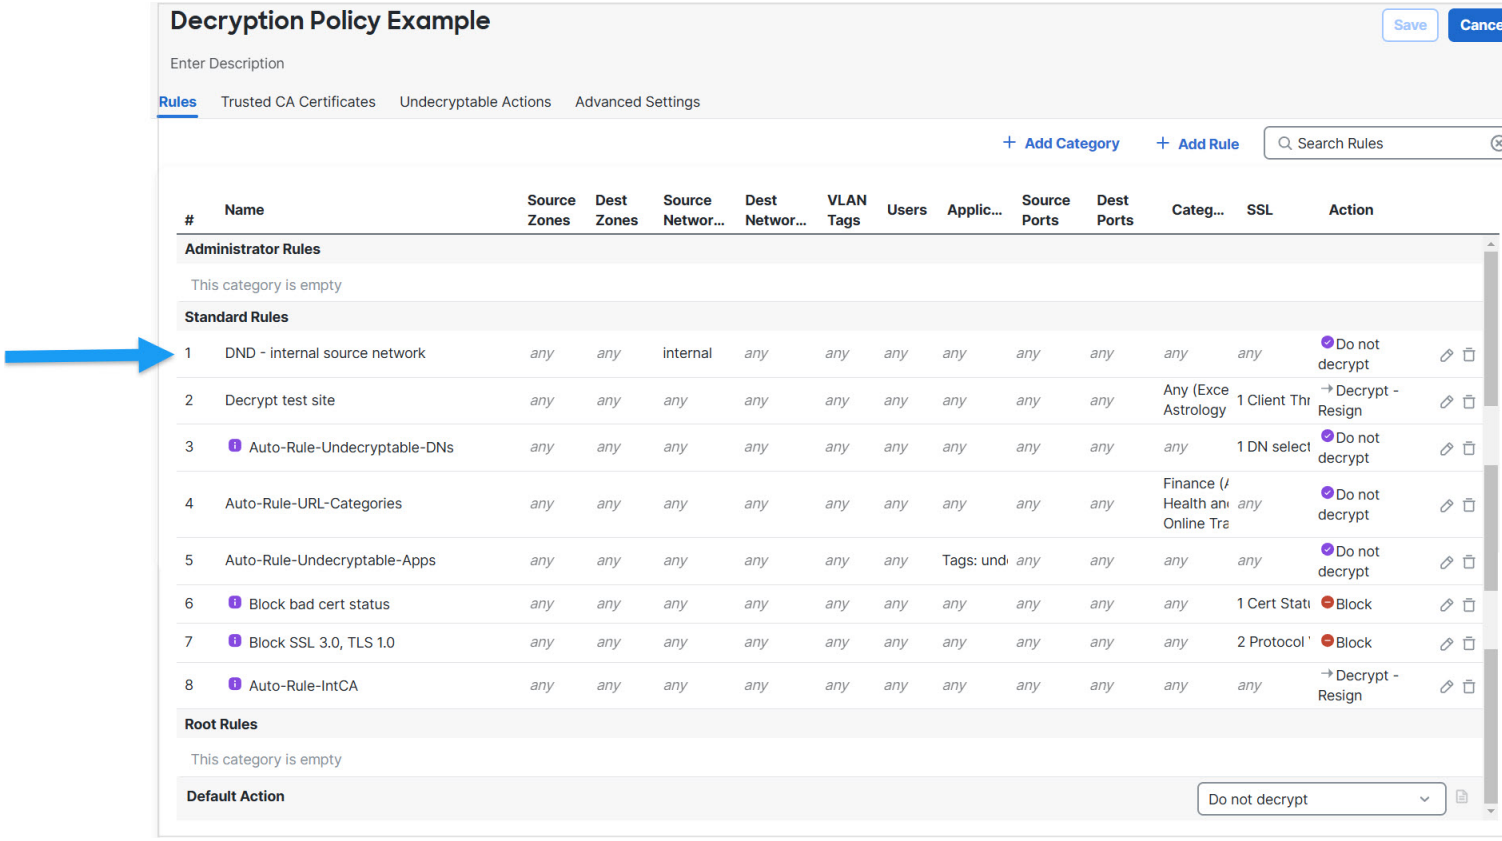

Note If you have traffic going from internal DNS servers to internal DNS resolvers (such as Cisco Umbrella Virtual Appliances), you can add **Do Not Decrypt** rules for them as well. You can even add those to prefiltering policies if the internal DNS servers do their own logging.

However, we strongly recommend you *do not* use **Do Not Decrypt** rules or prefiltering for DNS traffic that goes to the internet, such as internet root servers (for example, Microsoft internal DNS resolvers built into Active Directory). In those cases, you should fully inspect the traffic or even consider blocking it.

Rule detail:

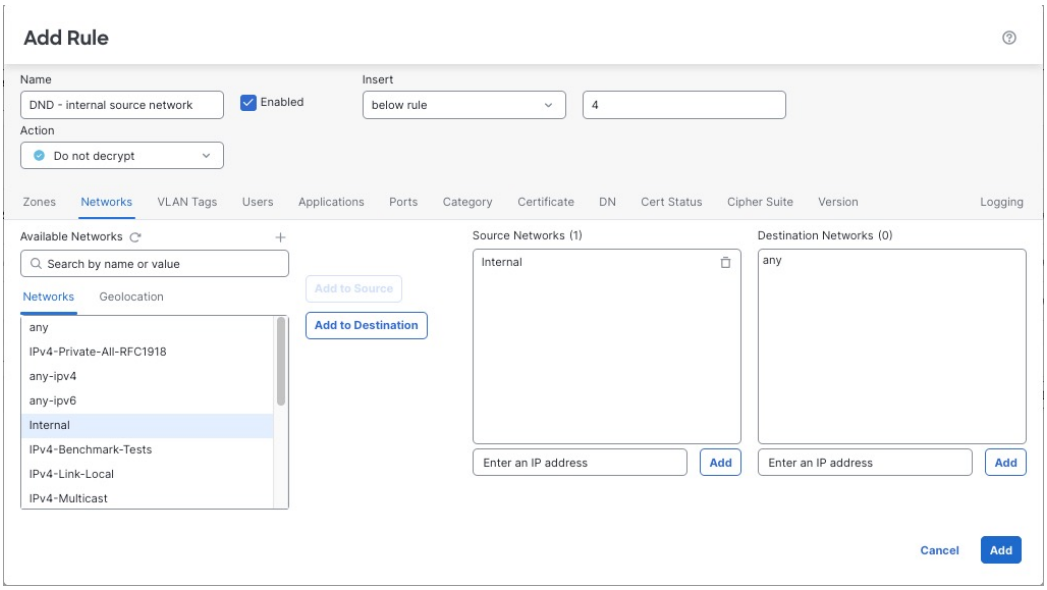

Next Manual Rule: Decrypt Specific Test Traffic

The next rule is *optional* in the example; use it to decrypt and monitor limited types of traffic before determining whether or not to allow it on your network.

After you run the decryption policy wizard, edit the policy to add the following rule. Drag it to the second position in the list of rules.

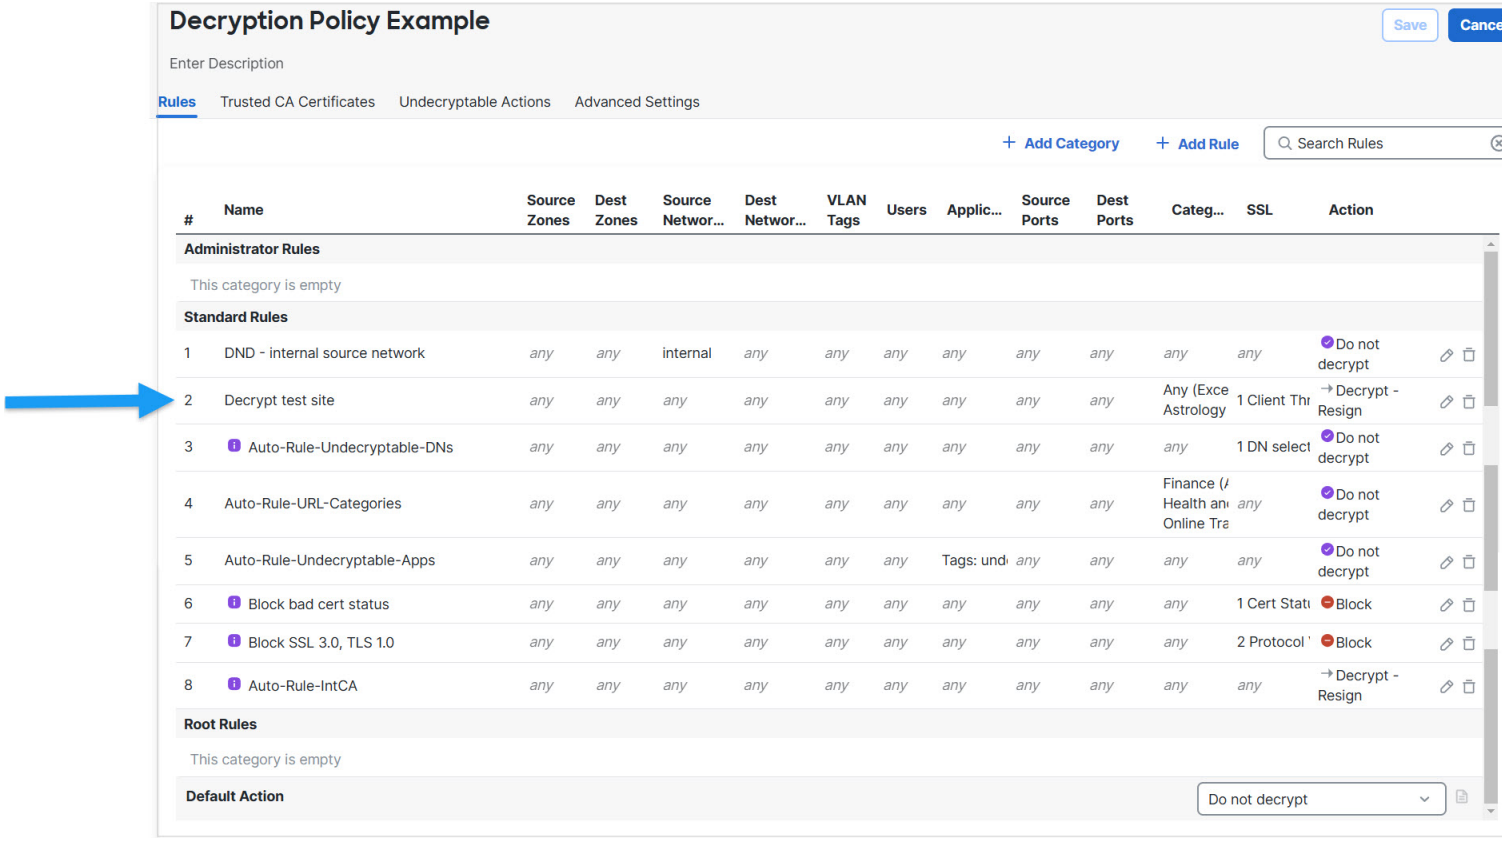

Rule detail:

Last Manual Decryption Rules: Block or Monitor Certificates and Protocol Versions

The last decryption rules, because they are the most specific and require the most processing, are rules that either monitor or block bad certificates and unsecure protocol versions.

After you run the decryption policy wizard, edit the policy to add the following rules. Drag them to a position *before* the **Decrypt - Resign** rule.

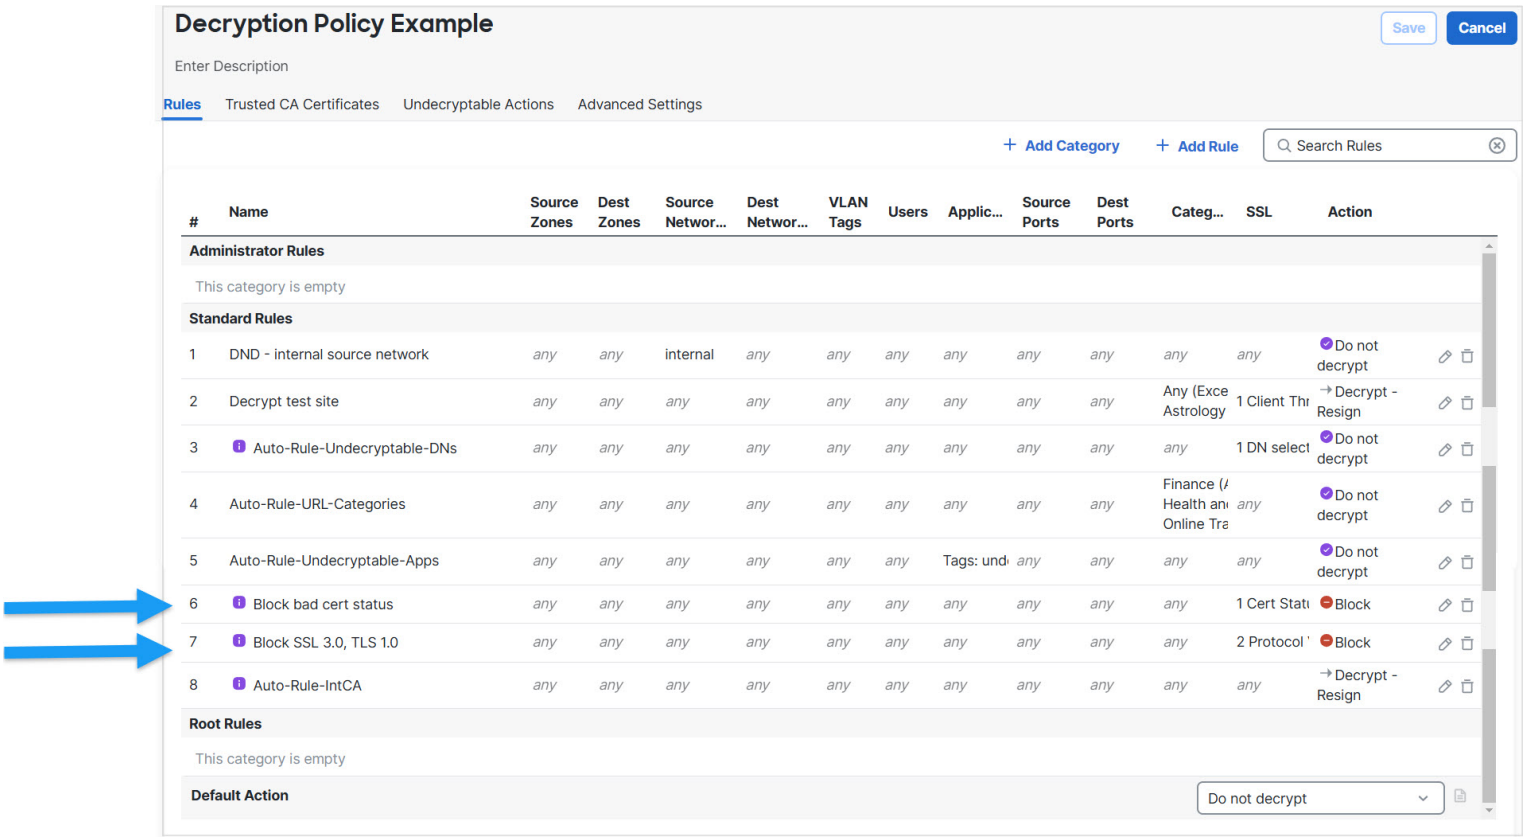

Rule details:

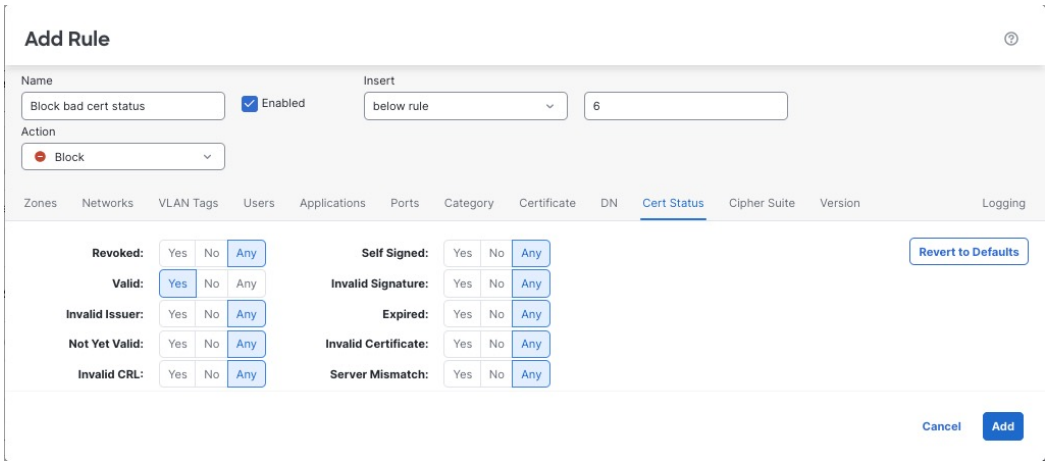

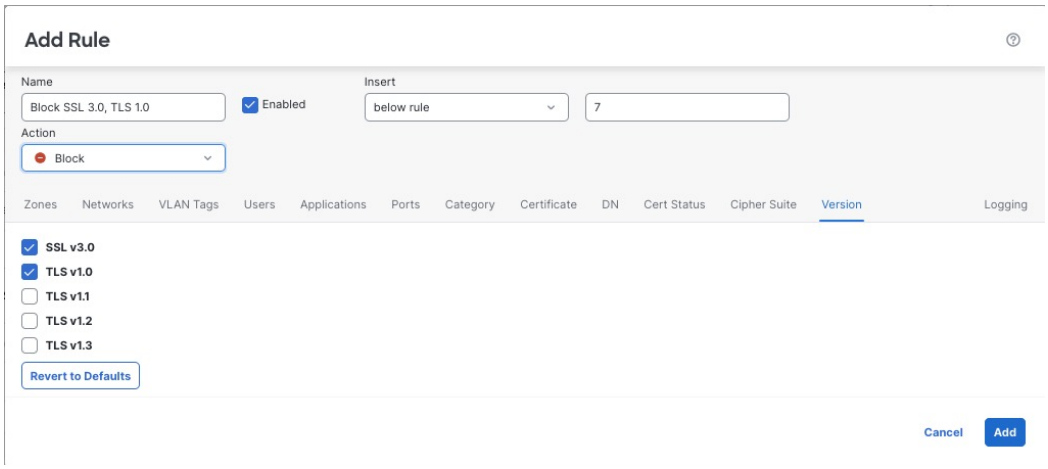

Example: Decryption Rule to Monitor or Block Certificate Status

The last decryption rules, because they are the most specific and require the most processing, are rules that either monitor or block bad certificates and unsecure protocol versions. The example in this section shows how to monitor or block traffic by certificate status.

Important Use the **Cipher Suite** and **Version** rule conditions *only* in rules with either the **Block** or **Block with reset** rule actions. Do not use **Cipher Suite** and **Version** with **Decrypt - Resign** or **Decrypt - Known Key** rule actions. These conditions in rules with other rule actions can interfere with the system's ClientHello processing, resulting in unpredictable performance.

Procedure

Step 1 Log in to the Secure Firewall Management Center if you haven't already done so.

Step 2 Click **Policies** > **Access Control heading** > **Decryption**.

Step 3 Click **Edit** (✎) next to your decryption policy.

Step 4 Click **Edit** (✎) next to a decryption rule.

Step 5 Click **Add Rule**.

Step 6 n the Add Rule dialog box, in the **Name** field, enter a name for the rule.

Step 7 Click **Cert Status**.

Step 8 For each certificate status, you have the following options:

- Click **Yes** to match against the *presence* of that certificate status.
- Click **No** to match against the *absence* of that certificate status.
- Click **Any** to skip the condition when matching the rule. In other words, choosing **Any** means the rule matches whether the certificate status is present or absent.

Step 9 From the **Action** list, click either **Monitor** to only monitor and log traffic that matches the rule or click **Block** or **Block with Reset** to block the traffic and optionally reset the connection.

Step 10 To save changes to the rule, at the bottom of the page, click **Add**.

Step 11 To save changes to the policy, at the top of the page, click **Save**.

Example

The organization trusts the Verified Authority certificate authority. The organization does not trust the Spammer Authority certificate authority. The system administrator uploads the Verified Authority certificate and an intermediate CA certificate issued by Verified Authority to the system. Because Verified Authority revoked a certificate it previously issued, the system administrator uploads the CRL that Verified Authority provided.

The following figure shows a certificate status rule condition checking for valid certificates, those issued by a Verified Authority, are not on the CRL, and still within the Valid From and Valid To date. Because of the configuration, traffic encrypted with these certificates is not decrypted and inspected with access control.

The following figure shows a certificate status rule condition checking for the absence of a status. In this case, because of the configuration, it matches against traffic encrypted with a certificate that has not expired.

In the following example, traffic would match this rule condition if the incoming traffic is using a certificate that has an invalid issuer, is self-signed, expired, and it is an invalid certificate.

The following graphic illustrates a certificate status rule condition that matches if the SNI of the request matches the server name or if the CRL is not valid.

Example: Decryption Rule to Monitor or Block Protocol Versions

This example shows how to block TLS and SSL protocols on your network that are no longer considered secure, such as TLS 1.0, TLS 1.1, and SSLv3. It's included to give you a little more detail about how protocol version rules work.

You should exclude nonsecure protocols from your network because they are all exploitable. In this example:

- You can block some protocols using **Version** page on the decryption rule.
- Because the system considers SSLv2 as undecryptable, you can block it using the **Undecryptable Actions** on the decryption policy.
- Similarly, because compressed TLS/SSL is not supported, you should block it as well.

Important Use the **Cipher Suite** and **Version** rule conditions *only* in rules with either the **Block** or **Block with reset** rule actions. Do not use **Cipher Suite** and **Version** with **Decrypt - Resign** or **Decrypt - Known Key** rule actions. These conditions in rules with other rule actions can interfere with the system's ClientHello processing, resulting in unpredictable performance.

Procedure

Step 1 Log in to the Secure Firewall Management Center if you haven't already done so.

Step 2 Click **Policies** > **Access Control heading** > **Decryption**.

Step 3 Click **Edit** (✎) next to your decryption policy.

Step 4 Click **Edit** (✎) next to a decryption rule.

Step 5 Click **Add Rule**.

Step 6 In the Add Rule dialog box, in the **Name** field, enter a name for the rule.

Step 7 From the **Action** list, click **Block** or **Block with reset**.

Step 8 Click **Version** page.

Step 9 Check the check boxes for protocols that are no longer secure, such as **SSL v3.0**, **TLS 1.0**, and **TLS 1.1**. Clear the check boxes for any protocols that are still considered secure.

The following figure shows an example.

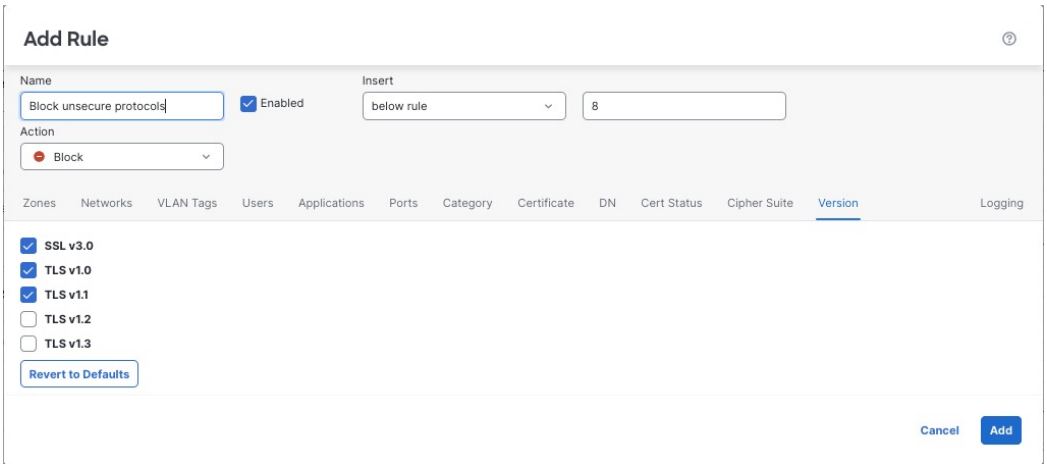

Step 10 Choose other rule conditions as needed.

Step 11 Click **Add**.

Optional Example: Manual Decryption Rule to Monitor or Block Certificate Distinguished Name

This rule is included to give you an idea about how to monitor or block traffic based on the server certificate's distinguishedname. It's included to give you a little more detail.

The distinguished name can consist of country code, common name, organization, and organizational unit, but typically consists of a common name only. For example, the common name in the certificate for `https://www.cisco.com` is `cisco.com`. (However, it's not always this simple; Distinguished Name (DN) Rule Conditions, on page 2323 shows how to find common names.)

The host name portion of the URL in the client request is the Server Name Indication (SNI). The client specifies which hostname they want to connect to (for example, `auth.amp.cisco.com`) using the SNI extension in the TLS handshake. The server then selects the corresponding private key and certificate chain that are required to establish the connection while hosting all certificates on a single IP address.

Procedure

Step 1 Log in to the Secure Firewall Management Center if you haven't already done so.

Step 2 Click **Policies** > **Access Control heading** > **Decryption**.

Step 3 Click **Edit** (✎) next to your decryption policy.

Step 4 Click **Edit** (✎) next to a decryption rule.

Step 5 Click **Add Rule**.

Step 6 In the Add Rule dialog box, in the **Name** field, enter a name for the rule.

Step 7 From the **Action** list, click **Block** or **Block with reset**.

Step 8 Click **DN**.

Step 9 Find the distinguished names you want to add from the **Available DNs**, as follows:

- To add a distinguished name object on the fly, which you can then add to the condition, click **Add** (+) above the **Available DNs** list.
- To search for distinguished name objects and groups to add, click the **Search by name or value** prompt above the **Available DNs** list, then type either the name of the object, or a value in the object. The list updates as you type to display matching objects.

Step 10 To select an object, click it. To select all objects, right-click and then **Select All**.

Step 11 Click **Add to Subject** or **Add to Issuer**.

Tip
You can also drag and drop selected objects.

Step 12 Add any literal common names or distinguished names that you want to specify manually. Click the **Enter DN or CN** prompt below the **Subject DNs** or **Issuer DNs** list; then type a common name or distinguished name and click **Add**.

Although you can add a CN or DN to either list, it's more common to add them to the **Subject DNs** list.

Step 13 Add or continue editing the rule.

Step 14 When you're done, to save changes to the rule, click **Add** at the bottom of the page.

Step 15 To save changes to the policy, click **Save** at the top of the page.

Example

The following figure shows a distinguished name rule condition searching for certificates issued to goodbakery.example.com or issued by goodca.example.com. Traffic encrypted with these certificates is allowed, subject to access control.

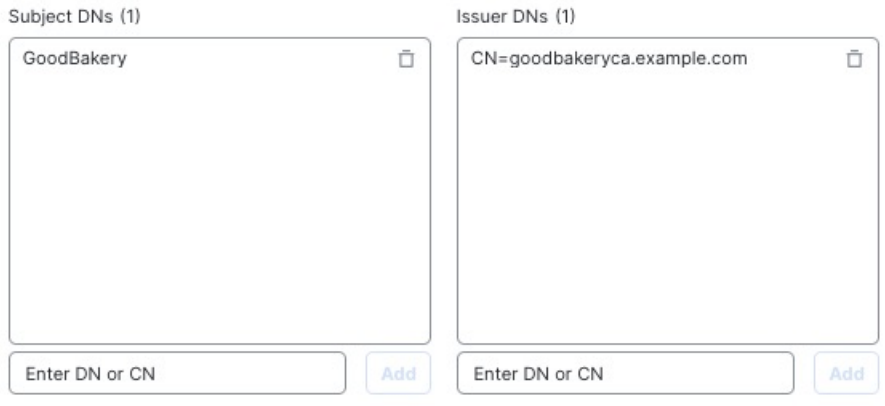

Associate the Decryption Policy with an Access Control Policy and Advanced Settings

This task discusses how to associate the decryption policy with an access control policy and setting recommended advanced settings for the access control policy.

For your decryption policy to be used by the system, you *must* associate it with an access control policy.

Before you begin

Create the sample decryption policy as discussed in this guide.

For more information about decryption policy advanced options, see Decryption Policy Advanced Options, on page 2298.

Procedure

Step 1 Log in to the Secure Firewall Management Center if you haven't already done so.

Step 2 Click **Policies** > **Access Control heading** > **Access Control**.

Step 3 Either create a new access control policy or click **Edit** (✎) to edit an existing one.

Step 4 Click the word Decryption as the following figure shows.

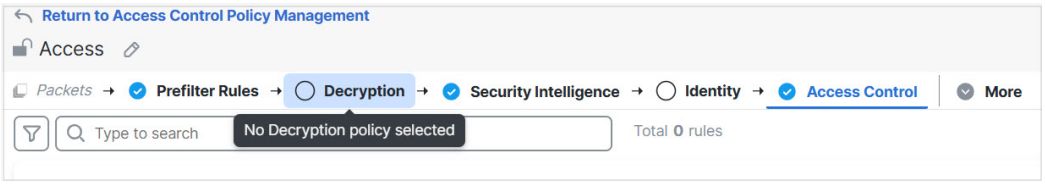

Step 5 From the list, click the name of your decryption policy and also check **Early application detection and URL categorization** as the following figure shows.

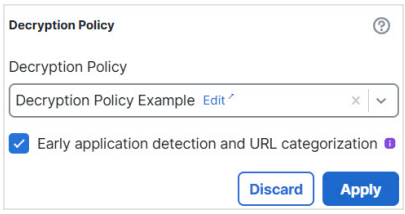

Step 6 Click **Apply**.

Step 7 Click **More** > **Advanced Settings** as the following figure shows.

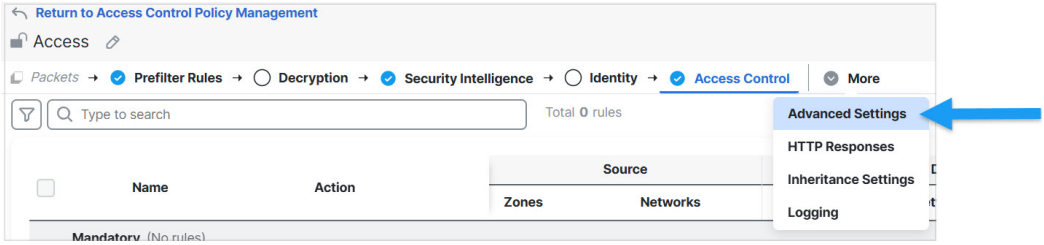

Step 8 Click **Edit** (✎) next to **TLS Server Identity Discovery**.

Step 9 Select the check box as the following figure shows.

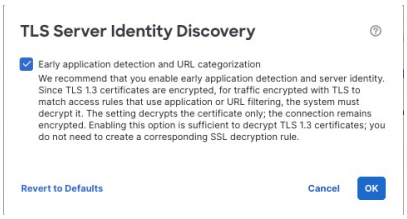

Step 10 Click **OK**.

Step 11 At the top of the page, click **Save**.

Step 12 At the top of the page, click **Return to Access Control Policy Management**, as the following figure shows

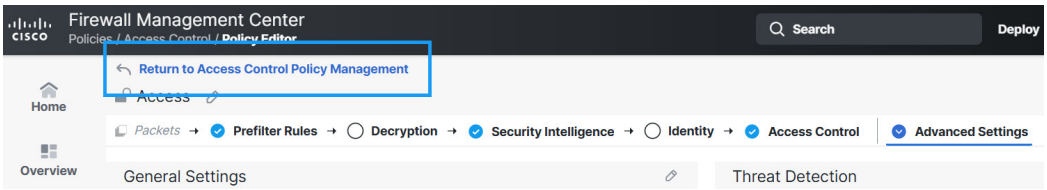

Step 13 Click **Edit** (✎) to edit the access control rule.

Step 14 At the bottom of the page, next to the default action, click (Default Logging and Inspection).

Step 15 Check **Log at beginning of connection** and any other options you choose.

For more information, see Logging Settings for Access Control Policies, on page 1755Logging Settings for Access Control Policies in the *Cisco Secure Firewall Management Center Device Configuration Guide*.

Step 16 Click **Apply**.

Step 17 At the top of the page, click **Save**.

What to do next

- Add rule conditions: Decryption Rule Conditions, on page 2316.
- Add a default policy action: Decryption Policy Default Actions, on page 2295.
- Configure logging options for the default action as described in *Logging Connections with a Policy Default Action* in the Cisco Secure Firewall Management Center Administration Guide .

- Set advanced policy properties: Decryption Policy Advanced Options, on page 2298.
- Associate the decryption policy with an access control policy as described in Associating Other Policies with Access Control, on page 1762.
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Traffic to Prefilter

Prefiltering is the first phase of access control, before the system performs more resource-intensive evaluation. Prefiltering is simple, fast, and early compared to subsequent evaluation, which uses inner headers and has more robust inspection capabilities.

Based on your security needs and traffic profile, you should consider prefiltering and therefore excluding from any policy and inspection the following:

- Common intraoffice applications such as Microsoft Outlook 365
- Elephant flows, such as server backups

Decryption Rule Settings

How to configure recommended best practice settings for your decryption rules.

Decryption rules: Enable logging for every rule except those with a **Do Not Decrypt** rule action. (It's up to you; if you want to see information about traffic that isn't decrypted, enable logging for those rules also.)

Procedure

Step 1 Log in to the Secure Firewall Management Center if you haven't already done so.

Step 2 Click **Policies** > **Access Control heading** > **Decryption**.

Step 3 Click **Edit** (✎) next to your decryption policy.

Step 4 Click **Edit** (✎) next to a decryption rule.

Step 5 Click the **Logging** tab.

Step 6 Click **Log at End of Connection**.

Step 7 Click **Save**.

Step 8 Click **Save** at the top of the page.

PART XII

User Identity

- User Identity Overview, on page 2371
- Realms, on page 2391
- User Control with the Passive Identity Agent, on page 2467
- User Control with ISE/ISE-PIC, on page 2499
- User Control with Captive Portal, on page 2525
- User Control with Remote Access VPN, on page 2545
- User Control with TS Agent, on page 2549
- User Identity Policies, on page 2553

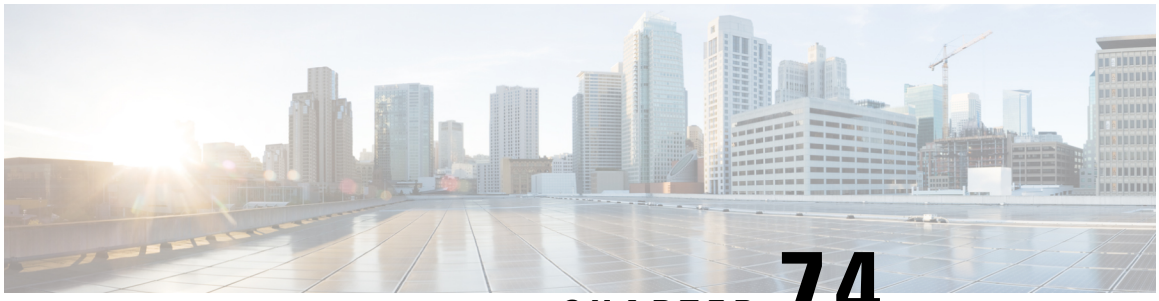

CHAPTER 74

User Identity Overview

The following topics discuss user identity:

- About User Identity, on page 2371
- Host and User Limits, on page 2383
- Identity Realm Limit, on page 2387
- User Limits for Microsoft Azure Active Directory Realms, on page 2388

About User Identity

User identity information can help you to identify the source of policy breaches, attacks, or network vulnerabilities, and trace them to specific users. For example, you could determine:

- Who owns the host targeted by an intrusion event that has a Vulnerable (level 1: red) impact level.
- Who initiated an internal attack or portscan.
- Who is attempting unauthorized access to a specified host.
- Who is consuming an unreasonable amount of bandwidth.
- Who has not applied critical operating system updates.
- Who is using instant messaging software or peer-to-peer file-sharing applications in violation of company policy.
- Who is associated with each indication of compromise on your network.

Armed with this information, you can use other features of the system to mitigate risk, perform access control, and take action to protect others from disruption. These capabilities also significantly improve audit controls and enhance regulatory compliance.

After you configure user identity sources to gather user data, you can perform user awareness and user control.

For more information about identity sources, see About User Identity Sources, on page 2372.

Related Topics

Identity Terminology, on page 2372
About User Identity Sources, on page 2372
Identity Deployments, on page 2376
How to Set Up an Identity Policy, on page 2378

Identity Terminology

This topic discusses common terminology for user identity and user control.

User awareness

Identifying users on your network using *identity sources* (such as or TS Agent). User awareness enables you to identify users from both *authoritative* (such as Active Directory) and *non-authoritative* (application-based) sources. To use Active Directory as an identity source, you must configure a realm and directory. For more information, see About User Identity Sources, on page 2372.

User control

Configuring an *identity policy* that you associate with an *access control policy*. (The identity policy is then referred to as an access control *subpolicy*.) The identity policy specifies the identity source and, optionally, users and groups belonging to that source.

By associating the identity policy with an access control policy, you determine whether to monitor, trust, block, or allow users or user activity in traffic on your network. For more information, see Access Control Policies, on page 1745.

Authoritative identity sources

A trusted server validated the user login (for example, Active Directory). You can use the data obtained from authoritative logins to perform user awareness and user control. Authoritative user logins are obtained from passive and active authentications:

- *Passive authentications* occur when a user authenticates with Cisco ISE.
- *Active authentications* occur when a user authenticates through preconfigured managed devices. Captive portal is another name for active authentication. Active authentication generally uses the same user repositories as passive authentication (the exceptions being ISE/ISE-PIC, TS Agent, and the Passive Identity Agent, which are passive only).

Non-authoritative identity sources

An unknown or untrusted server validated the user login. Traffic-based detection is the only non-authoritative identity source supported by the system. You can use the data obtained from non-authoritative logins to perform user awareness.

About User Identity Sources

The following table provides a brief overview of the user identity sources supported by the system. Each identity source provides a store of users for user awareness. These users can then be controlled with identity and access control policies.

User Identity Source	Server Requirements	Login Type	Authentication Type	User Control	For more, see...
Captive portal	OpenLDAP Microsoft Active Directory Microsoft Azure Active Directory	Authoritative	Active	Yes	The Captive Portal Identity Source, on page 2525 Create a Microsoft Azure AD (SAML) Realm for Active Authentication (Captive Portal), on page 2412
Passive authentication	OpenLDAP Microsoft Active Directory	Non-authoritative	Active	Yes	Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426
Passive authentication	Microsoft Azure Active Directory	Passive	Passive	Yes	Configure Microsoft Entra ID for Passive Authentication, on page 2395
Passive authentication with the Passive Identity Agent	Microsoft Active Directory	Non-authoritative	Passive	Yes	The Passive Identity Agent Identity Source, on page 2467
Passive authentication with the TS Agent	Microsoft Windows Terminal Server	Authoritative	Passive	Yes	The Terminal Services (TS) Agent Identity Source, on page 2549
Remote Access VPN	OpenLDAP or Microsoft Active Directory	Authoritative	Active	Yes	The Remote Access VPN Identity Source, on page 2545
	RADIUS	Authoritative	Active	No, awareness only	
ISE/ISE-PIC	Microsoft Active Directory	Authoritative	Passive	Yes	The ISE/ISE-PIC Identity Source, on page 2499
Traffic-based detection (Configured in the network discovery policy.)	—	Non-authoritative	—	No, awareness only	The Traffic-Based Detection Identity Source, on page 2651

Consider the following when selecting identity sources to deploy:

- You must use traffic-based detection for non-LDAP user logins.
- You must use traffic-based detection or captive portal to record failed login or authentication activity. A failed login or authentication attempt does not add a new user to the list of users in the database.
- The captive portal identity source requires a managed device with a routed interface. You *cannot* use an inline (also referred to as tap mode) interface with captive portal.

Data from those identity sources is stored in the Secure Firewall Management Center's users database and the user activity database. You can configure Secure Firewall Management Center-server user downloads to automatically and regularly download new user data to your databases.

After you configure identity rules using the desired identity source, you must associate each rule with an access control policy and deploy the policy to managed devices for the policy to have any effect. For more information about access control policies and deployment, see Associating Other Policies with Access Control, on page 1762.

For general information about user identity, see About User Identity, on page 2371.

Best Practices for User Identity

We recommend you review the following information before you set up identity policies.

- Know user limits
- Health monitor
- Use latest version of ISE/ISE-PIC, two types of remediation
- User agent support drops in 6.7
- Captive portal requires routed interface, several individual tasks

Microsoft Active Directory and LDAP

The system supports Active Directory, LDAP, and other user repositories for user awareness and control. The association between an Active Directory or LDAP repository and the Secure Firewall Management Center is referred to as a *realm*. You should create one realm per LDAP server or Active Directory domain. For details about which versions are supported, see Supported Servers for Realms, on page 2434.

The only user identity source supported by LDAP is captive portal. To use other identity sources (with the exception of ISE/ISE-PIC), you must use Active Directory.

For Active Directory only:

- Create one *directory* per domain controller.

  For details, see Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426

- Users and groups in trust relationships between two domains are supported provided you add all Active Directory domains and domain controllers as realms and directories, respectively.

  For more information, see Realms and Trusted Domains, on page 2431.

Passive Identity Agent identity source

The passive identity agent software can monitor multiple Microsoft Active Directory servers and domain controllers and send user name and IP address information to the Secure Firewall Management Center.

For more information, see The Passive Identity Agent Identity Source, on page 2467.

Health monitor

The Secure Firewall Management Center health monitor provides valuable information about the status of various Secure Firewall Management Center functions, including:

- User/realm mismatches
- Snort memory usage
- ISE connection status

For more information about health modules, see *Health Modules* in the Cisco Secure Firewall Management Center Administration Guide.

To set up policies to monitor health modules, see *Creating Health Policies* in the Cisco Secure Firewall Management Center Administration Guide.

Device-specific user limits

Every physical or virtual Firewall Management Center device has limits to the number of users that can be downloaded. When the user limit is reached, the Firewall Management Center can run out of memory and can function unreliably as a result.

User limits are discussed in User Limits for Microsoft Active Directory, on page 2384.

If you use the ISE/ISE-PIC identity source, you can optionally limit the subnets the Firewall Management Center monitors to reduce memory usage using identity mapping filters as discussed in Create an Identity Policy, on page 2555.

Use the latest version of ISE/ISE-PIC

If you expect to use the ISE/ISE-PIC identity source, we strongly recommend you always use the latest version to make sure you get the latest features and bug fixes.

pxGrid 2.0 (which is used by version 2.6 patch 6 or later; or 2.7 patch 2 or later) also changes the remediation used by ISE/ISE-PIC from Endpoint Protection Service (EPS) to Adaptive Network Control (ANC). If you upgrade ISE/ISE-PIC, you must migrate your mediation policies from EPS to ANC.

More information about using ISE/ISE-PIC can be found in ISE/ISE-PIC Guidelines and Limitations, on page 2501.

To set up the ISE/ISE-PIC identity source, see How to Configure ISE/ISE-PIC for User Control, on page 2504.

Captive portal information

You can use captive portal active authentication with any of the following:

- LDAP

  For more information, see The Captive Portal Identity Source, on page 2525.

- Microsoft AD

For more information, see The Captive Portal Identity Source, on page 2525.

- Microsoft Azure AD (SAML)

  For more information, see Create a Microsoft Azure AD (SAML) Realm, on page 2392.

For more information about the captive portal identity source, see The Captive Portal Identity Source, on page 2525.

TS Agent information

The TS Agent user identity source is required to identify user sessions on a Windows Terminal Server. The TS Agent software must be installed on the Terminal Server machine as discussed in the *Cisco Terminal Services (TS) Agent Guide*. In addition, you must synchronize the time on your TS Agent server with the time on the Secure Firewall Management Center.

TS Agent data is visible in the Users, User Activity, and Connection Event tables and can be used for user awareness and user control.

For more information, see TS Agent Guidelines, on page 2550.

Associate the identity policy with an access control policy

After you configure your realm, directory, and user identity source, you must set up identity rules in an identity policy. To make the policy effective, you must associate the identity policy with an access control policy.

For more information about creating an identity policy, see Create an Identity Policy, on page 2555.

For more information about creating identity rules, see Create an Identity Rule, on page 2563.

To associate an identity policy with an access control policy, see Associating Other Policies with Access Control, on page 1762.

Identity Deployments

When the system detects user data from a user login, from any identity source, the user from the login is checked against the list of users in the Secure Firewall Management Center user database. If the login user matches an existing user, the data from the login is assigned to the user. Logins that do not match existing users cause a new user to be created, unless the login is in SMTP traffic. Non-matching logins in SMTP traffic are discarded.

The group to which the user belongs is associated with the user as soon as the user is seen by the Secure Firewall Management Center.

The following diagram illustrates how the system collects and stores user data:

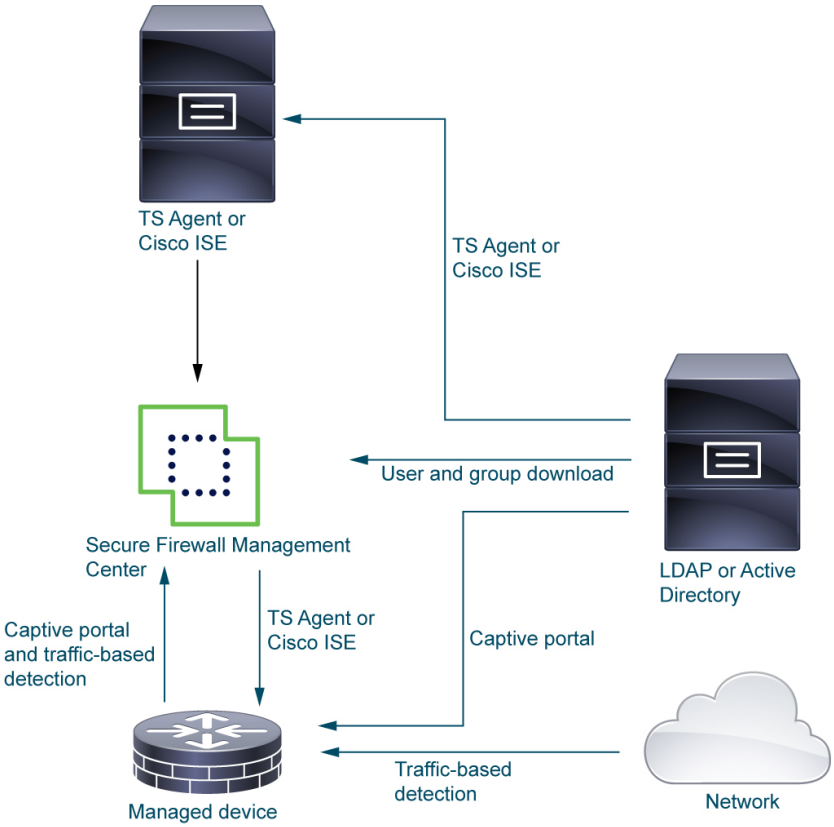

Sample identity deployments

In the preceding figure, Secure Firewall Management Center and one managed device are deployed to AWS and the other devices are located on premises. These devices can be either physical or virtual; they just need to be able to communicate with each other.

LDAP or Active Directory are needed only for TS Agent and captive portal, as the following paragraphs explain.

For more information about setting up a system like this, see How to Set Up an Identity Policy, on page 2378.

ISE/ISE-PIC identity source

When you deploy the ISE/ISE-PIC identity source, the ISE/ISE-PIC server sends users, groups, and subscriptions to the Secure Firewall Management Center.

You can optionally have an LDAP server in an ISE/ISE-PIC deployment but because it's optional, it isn't shown in the following figure.

For more information about ISE/ISE-PIC, see The ISE/ISE-PIC Identity Source, on page 2499.

Passive identity agent identity source

Software that monitors one or several Microsoft AD domain controllers, sending user name and IP address information to the Secure Firewall Management Center.

For more information, see The Passive Identity Agent Identity Source, on page 2467.

TS Agent identity source

The Terminal Services (TS) Agent software runs on a Microsoft Server and sends Secure Firewall Management Center user information based on the port range the users log in to the server with. TS Agent gets user identity information from LDAP or Active Directory and sends it to Secure Firewall Management Center.

For more information about the TS Agent identity source, see The Terminal Services (TS) Agent Identity Source, on page 2549.

Captive portal

You can use captive portal active authentication with any of the following:

- LDAP

  For more information, see The Captive Portal Identity Source, on page 2525.

- Microsoft AD

  For more information, see The Captive Portal Identity Source, on page 2525.

- Microsoft Azure AD (SAML)

  For more information, see Create a Microsoft Azure AD (SAML) Realm, on page 2392.

For more information about the captive portal identity source, see The Captive Portal Identity Source, on page 2525.

How to Set Up an Identity Policy

This topic provides a high-level overview of setting up an identity policy using any available user identity source: TS Agent, ISE/ISE-PIC, Microsoft Azure AD (SAML), captive portal.

Procedure

	Command or Action	Purpose
Step 1	Choose your identity source.	• To use Microsoft Active Directory or LDAP, continue with the next step. • To use a Microsoft Azure AD (SAML) realm, see Create a Microsoft Azure AD (SAML) Realm, on page 2392. • To use the Cisco ISE identity source, see Ways to Configure the Cisco Identity Services Engine (Cisco ISE) Identity Source, on page 2512. • To use the passive identity agent identity source, see The Passive Identity Agent Identity Source, on page 2467.
Step 2	(Optional.) Create a realm and directory, one realm for every domain in the forest that contain users you want to use in user control.	Creating a realm, realm directory is optional if any of the following are true:

	Command or Action	Purpose
	Also create one directory for every domain controller. Only users and groups that have corresponding Firewall Management Center realms and directories can be used in identity policies..	• You use SGT ISE attribute conditions but not user, group, realm, Endpoint Location, or Endpoint Profile conditions. • You are using an identity policy only to filter network traffic. The *realm* is a trusted user and group store, typically a Microsoft Active Directory repository. The Firewall Management Center downloads users and groups at intervals you specify. You can include or exclude users and groups from being downloaded. See Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426. For details about the options to create a realm, see Realm Fields, on page 2436. A *directory* is an Active Directory domain controller that organizes information about a computer network's users and network shares. An Active Directory controller provides Directory Services for the realm. Active Directory distributes user and group objects across multiple domain controllers, which are peers that propagate local changes between each other by the use of Directory Services. For more information, see the Active Directory technical specification glossary on MSDN. You can specify more than one directory for a realm, in which case each domain controller is queried in the order listed on the realm's **Directory** tab page to match user and group credentials for user control. **Note** Configuring a realm or realm sequence is optional if you plan to configure SGT ISE attribute conditions but not user, group, realm, Endpoint Location, or Endpoint Profile conditions.
Step 3	Synchronize users and groups from the realm.	To be able to control users and groups, you must synchronize them with the Firewall Management Center. You can synchronize them with users and groups whenever you want or you can configure the system to synchronize them at a specified interval.

	Command or Action	Purpose
		When you synchronize users and groups, you can specify exceptions; for example, you can exclude the Engineering group from all user control for that realm, or you can exclude the user **joe.smith** from user controls that apply to the Engineering group. See Synchronize Users and Groups, on page 2445
Step 4	(Optional.) Create a realm sequence.	A realm sequence is an ordered list of realms that, when used in an identity policy, causes the system to search the realms in the specified order to find users to match the rule. See Create a Realm Sequence, on page 2446.
Step 5	Create a method to retrieve user and group data (the *identity source*).	Set up an identity source with its unique configuration to be able to control users and groups using data stored in the realm. Identity sources include TS Agent, captive portal, or Remote VPN. See one of the following: • How to Configure the Captive Portal for User Control, on page 2529 • Ways to Configure the Cisco Identity Services Engine (Cisco ISE) Identity Source, on page 2512 • Configure RA VPN for User Control, on page 2546
Step 6	Create an identity policy.	An identity policy contains one or more identity rules, optionally organized in categories. See Create an Identity Policy, on page 2555. **Note** Configuring a realm or realm sequence is optional if you plan to configure SGT ISE attribute conditions but not user, group, realm, Endpoint Location, or Endpoint Profile conditions; or if you use your identity policy only to filter network traffic.
Step 7	Create one or more identity rules.	Identity rules enable you to specify a number of matching criteria, including the type of authentication, network zones, networks or geolocation, realms, realm sequences, and so on. See Create an Identity Rule, on page 2563.
Step 8	Associate your identity policy with an access control policy.	An access control policy filters and optionally inspects traffic. An identity policy must be

	Command or Action	Purpose
		associated with an access control policy to have any effect. See Associating Other Policies with Access Control, on page 1762.
Step 9	Deploy the access control policy to at least one managed device.	To use your policy to control user activity, the policy must be deployed to the managed devices to which clients connect. See Deploy Configuration Changes, on page 259.
Step 10	Monitor user activity.	View a list of active sessions collected by user identity sources or a list of user information collected by user identity sources. See *Using Workflows* in the Cisco Secure Firewall Management Center Administration Guide . An identity policy is not required if all of the following are true: • You use the ISE/ISE-PIC identity source. • You do not use users or groups in access control policies. • You use Security Group Tags (SGT) in access control policies. For more information, see ISE SGT vs Custom SGT Rule Conditions.

Related Topics

Configuring Traffic-Based User Detection, on page 2653

The User Activity Database

The user activity database on the Secure Firewall Management Center contains records of user activity on your network detected or reported by all of your configured identity sources. The system logs events in the following circumstances:

- When it detects individual logins or logoffs.
- When it detects a new user.
- When a system administrator manually delete a user.
- When the system detects a user that is not in the database, but cannot add the user because you have reached your user limit.
- When you resolve an indication of compromise associated with a user, or enable or disable indication of compromise rules for a user.

Note If the TS Agent monitors the same users as another passive authentication identity source (such as the ISE/ISE-PIC), the Firewall Management Center prioritizes the TS Agent data. If the TS Agent and another passive source report identical activity from the same IP address, only the TS Agent data is logged to the Firewall Management Center.

You can view user activity detected by the system using the Secure Firewall Management Center. (**Analysis** > **Users heading** > **User Activity**.)

The Users Database

The users database on the Secure Firewall Management Center contains a record for each user detected or reported by all of your configured identity sources. You can use data obtained from an authoritative source for user control.

See About User Identity Sources, on page 2372 for more information about the supported non-authoritative and authoritative identity sources.

The total number of users the Secure Firewall Management Center can store depends on the Secure Firewall Management Center model, as described in User Limits for Microsoft Active Directory, on page 2384. After the user limit is reached, the system prioritizes previously-undetected user data based on its identity source, as follows:

- If the new user is from a non-authoritative identity source, the system does not add the user to the database. To allow new users to be added, you must delete users manually or with a database purge.
- If the new user is from an authoritative identity source, the system deletes the non-authoritative user who has remained inactive for the longest period and adds the new user to the database.

If an identity source is configured to exclude specific user names, user activity data for those user names are not reported to the Secure Firewall Management Center. These excluded user names remain in the database, but are not associated with IP addresses. For more information about the type of data stored by the system, see *User Data* in the Cisco Secure Firewall Management Center Administration Guide.

If you have Firewall Management Center high availability configured and the primary fails, no logins reported by a captive portal, ISE/ISE-PIC, TS Agent, or Remote Access VPN device can be identified during failover downtime, even if the users were previously seen and downloaded to the Firewall Management Center. The unidentified users are logged as Unknown users on the Firewall Management Center. After the downtime, the Unknown users are reidentified and processed according to the rules in your identity policy.

Note If the TS Agent monitors the same users as another passive authentication identity source (ISE/ISE-PIC), the Firewall Management Center prioritizes the TS Agent data. If the TS Agent and another passive source report identical activity from the same IP address, only the TS Agent data is logged to the Firewall Management Center.

When the system detects a new user session, the user session data remains in the users database until one of the following occurs:

- A user on the Firewall Management Center manually deletes the user session.
- An identity source reports the logoff of that user session.

- A realm ends the user session as specified by the realm's **User Session Timeout: Authenticated Users**, **User Session Timeout: Failed Authentication Users**, or **User Session Timeout: Guest Users** setting.

Host and User Limits

Your Secure Firewall Management Center model determines how many individual hosts you can monitor with your deployment, as well as how many users you can monitor and use to perform user control.

Host Limit

The system adds a host to the network map when it detects activity associated with an IP address in your monitored network (as defined in your network discovery policy). The number of hosts a Secure Firewall Management Center can monitor, and therefore store in the network map, depends on its model.

Table 128: Host Limits by Secure Firewall Management Center Model

Firewall Management Center Model	Hosts
MC1000	50,000
MC1600	50,000
MC1700	50,000
MC2500	150,000
MC2600	150,000
MC2700	150,000
MC4500	600,000
MC4600	600,000
MC4700	600,000
virtual	50,000

You cannot view contextual data for hosts not in the network map. However, you can perform access control. For example, you can perform application control on traffic to and from a host not in the network map, even though you cannot use a compliance allow list to monitor the host's network compliance.

Note The system counts MAC-only hosts separately from hosts identified by both IP addresses and MAC addresses. All IP addresses associated with a host are counted together as one host.

Reaching the Host Limit and Deleting Hosts

The network discovery policy controls what happens when you detect a new host after you reach the host limit; you can drop the new host, or replace the host that has been inactive for the longest time. You can also set the period after which the system removes a host from the network map due to inactivity. Although you can manually delete a host, an entire subnet, or all of your hosts from the network map, if the system detects activity associated with a deleted host, it re-adds the host.

Related Topics

Network Discovery Data Storage Settings, on page 2659

User Limits for Microsoft Active Directory

About user limits

Your Firewall Management Center model determines how many individual users you can monitor. The user is added to the Firewall Management Center user database when:

- The user is downloaded from a realm.
- A captive portal or RA-VPN user logs in.
- A user is detected from any identity source (for example, TS Agent).

Only authoritative users are available for user control with access control policies.

Note the following:

- The maximum number of *downloaded* users depends on your Firewall Management Center model.
- The maximum number of *concurrent* user sessions (that is, logins) depends on your managed device model. A single user can have multiple sessions from different unique IP addresses.

Note The system downloads all user sessions to all Firewall Threat Defense devices. If you have devices with different user concurrent user session limits, the Firewall Threat Defense with the smallest limit reports health warnings when its memory reaches the configured limit. (For example, if your Firewall Management Center manages a Firepower 2110 and a 4125, the 2110 reports health warnings when the number of concurrent user sessions approaches its maximum of 64,000.)

User Limits for Microsoft Active Directory

Table 129: Maximum Concurrent User Login Limits by Firewall Threat Defense

Firewall Threat Defense Model	Maximum Concurrent User Logins per Realm
Firewall Threat Defense Virtual 5, 10, 20, 30, 50 (any supported hypervisor)	64,000

Firewall Threat Defense Model	Maximum Concurrent User Logins per Realm
Firepower 1010, 1120, 1140, 1150 Secure Firewall 1210, 1220, 1230, 1240, 1250 Firepower 2110, 2120, 2130 Secure Firewall 3105, 3110, 3120	64,000
Firepower 2140 Secure Firewall 3130, 3140 Firepower 4112, 4115, 4125	150,000
Firepower 4145 Firepower 9300	300,000
Secure Firewall 4215	300,000
Secure Firewall 4225, 4245	315,000

User limits are applied per Microsoft Active Directory realm. That is, if you attempt to download more than the maximum users in a *single* realm, the download stops after that many users and a health alert is displayed. If, however, you attempt to download more than the maximum number of users spread across *different* realms, the download succeeds (unless any one realm has more than 150,000 users, in which case the download fails for that realm).

Table 130: Maximum Downloaded Users by Firewall Management Center Model

Firewall Management Center Model	Maximum Downloaded Users
FMC 1000	50,000
FMC 1600	50,000
FMC 1700	50,000
FMC 2500	150,000
FMC 2600	150,000
FMC 2700	150,000
FMC 4500	600,000
FMC 4600	600,000
FMC 4700	600,000
Firewall Management Center Virtual (any supported hypervisor)	50,000
Firewall Management Center Virtual 300 (any supported hypervisor)	150,000

When the system detects a new, previously-undetected user after the limit has been reached, it prioritizes user data based on their identity source:

- If the new user is from a non-authoritative source, the system does not add the non-authoritative user to the database. To allow new users to be added, you must delete users manually or purge the database.
- If the new user is from an authoritative identity source, the system deletes the non-authoritative user who has remained inactive for the longest period of and adds the new authoritative user to the database.

  If there are only authoritative users, the system deletes the authoritative user who has remained inactive for the longest period of time and adds the new user to the database.

Troubleshooting information can be found in Troubleshoot User Control, on page 2573.

Tip Note that if you are using traffic-based detection, you can restrict user logging by protocol to help minimize username clutter and preserve space in the database. For example, you could prevent the system from adding users discovered in AIM, POP3, and IMAP traffic because you know it is traffic from specific contractors or visitors you do not want to monitor.

User Limits for Microsoft Azure Active Directory Realms

Microsoft Azure Active Directory User Limits

About user limits

Your Firewall Management Center model determines how many individual users you can monitor.

Note the following:

- The maximum number of *downloaded* users depends on your Firewall Management Center model.
- The maximum number of *concurrent* user sessions (that is, logins) depends on your managed device model. A single user can have multiple sessions from different unique IP addresses.

Note The system downloads all user sessions to all Firewall Threat Defense devices. If you have devices with different user concurrent user session limits, the Firewall Threat Defense with the smallest limit reports health warnings when its memory reaches the configured limit. (For example, if your Firewall Management Center manages a Firepower 2110 and a 4125, the 2110 reports health warnings when the number of concurrent user sessions approaches its maximum of 64,000.)

Refer to the following tables.

Table 131: Maximum Downloaded Users by Firewall Management Center Model[1]

Firewall Management Center Model	Number of Cisco Secure Dynamic Attributes Connector Connectors	Maximum Downloaded Users
FMC1600, FMC1700	10	50,000
FMC2600, FMC2700	20	150,000
FMC4600, FMC4700	30	600,000
Firewall Management Center Virtual (any supported hypervisor)	10	50,000
Firewall Management Center Virtual 300 (any supported hypervisor)	20	150,000

[1]—Firewall Management Center models are subject to end of life and end of sale. For more information, see End-Of-Life and End-Of-Sale Notices.

Table 132: Maximum Concurrent User Login Limits by Firewall Threat Defense

Firewall Threat Defense Model	Maximum Concurrent User Logins
Firewall Threat Defense Virtual 5, 10, 20, 30, 50 (any supported hypervisor)	50,000
Firepower 1010, 1120, 1140, 1150 Firepower 2110, 2120, 2130 Secure Firewall 3110, 3120, 3130, 3140	50,000
Firepower 2140 Secure Firewall 3130, 3140	150,000
Firepower 4140, 4145, 4150 Firepower 9300 Secure Firewall 4215	225,000
Secure Firewall 4225, 4245	300,000

Identity Realm Limit

We limit you to a maximum of 25 realms regardless of Firewall Threat Defense model; the same limit applies to the Cloud-Delivered Firewall Management Center.

The number of *users* you can download to a system is limited according to the information provided in:

User Limits for Microsoft Active Directory, on page 2384

User Limits for Microsoft Azure Active Directory Realms

Microsoft Azure Active Directory User Limits

About user limits

Your Firewall Management Center model determines how many individual users you can monitor.

Note the following:

- The maximum number of *downloaded* users depends on your Firewall Management Center model.
- The maximum number of *concurrent* user sessions (that is, logins) depends on your managed device model. A single user can have multiple sessions from different unique IP addresses.

Note The system downloads all user sessions to all Firewall Threat Defense devices. If you have devices with different user concurrent user session limits, the Firewall Threat Defense with the smallest limit reports health warnings when its memory reaches the configured limit. (For example, if your Firewall Management Center manages a Firepower 2110 and a 4125, the 2110 reports health warnings when the number of concurrent user sessions approaches its maximum of 64,000.)

Refer to the following tables.

Table 133: Maximum Downloaded Users by Firewall Management Center Model[1]

Firewall Management Center Model	Number of Cisco Secure Dynamic Attributes Connector Connectors	Maximum Downloaded Users
FMC1600, FMC1700	10	50,000
FMC2600, FMC2700	20	150,000
FMC4600, FMC4700	30	600,000
Firewall Management Center Virtual (any supported hypervisor)	10	50,000
Firewall Management Center Virtual 300 (any supported hypervisor)	20	150,000

[1]—Firewall Management Center models are subject to end of life and end of sale. For more information, see End-Of-Life and End-Of-Sale Notices.

Table 134: Maximum Concurrent User Login Limits by Firewall Threat Defense

Firewall Threat Defense Model	Maximum Concurrent User Logins
Firewall Threat Defense Virtual 5, 10, 20, 30, 50 (any supported hypervisor)	50,000
Firepower 1010, 1120, 1140, 1150 Firepower 2110, 2120, 2130 Secure Firewall 3110, 3120, 3130, 3140	50,000
Firepower 2140 Secure Firewall 3130, 3140	150,000
Firepower 4140, 4145, 4150 Firepower 9300 Secure Firewall 4215	225,000
Secure Firewall 4225, 4245	300,000

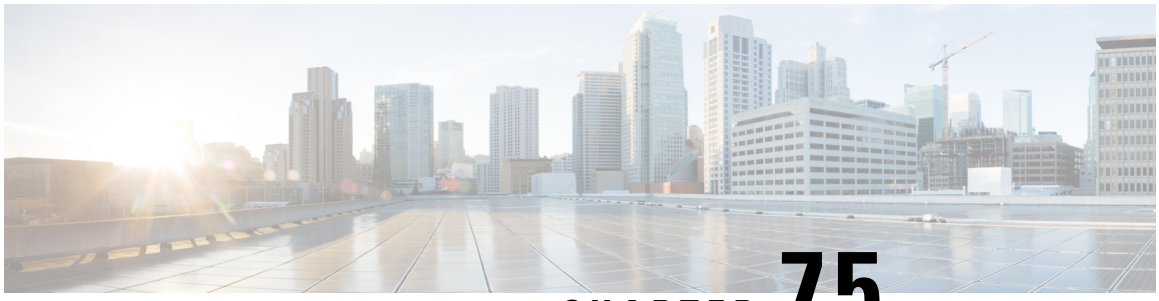

CHAPTER 75

Realms

The following topics describe realms and identity policies:

- License Requirements for Realms, on page 2391
- Requirements and Prerequisites for Realms, on page 2391
- Create a Microsoft Azure AD (SAML) Realm, on page 2392
- Create a Microsoft Azure AD (SAML) Realm for Passive Authentication, on page 2417
- Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426
- Create a Realm Sequence, on page 2446
- Configure the Firewall Management Center for Cross-Domain-Trust: The Setup, on page 2447
- Manage a Realm, on page 2455
- Compare Realms, on page 2456
- Troubleshoot Realms and User Downloads, on page 2456
- History for Realms, on page 2464

License Requirements for Realms

Threat Defense License

Any

Requirements and Prerequisites for Realms

Model Support

Any.

Supported Domains

Any

User Roles

- Admin

- Access Admin
- Network Admin

Create a Microsoft Azure AD (SAML) Realm

You can use a Microsoft Azure Active Directory (AD) (now called *Entra ID*) realm for either passive authentication or active authentication.

If you enabled Change Management, you must approve all certificates used in this procedure. Open a new ticket or edit an existing one. For more information, see Creating Change Management Tickets, on page 240 and .Policies and Objects that Support Change Management, on page 235

Passive authentication

Passive authentications occur when a user authenticates with Cisco ISE.

You have the following options, depending on your choice of user and group repository:

- To use Cisco ISE as a repository for users and to perform passive authentication using Entra ID. For more information, see:
  - About Entra ID and Cisco ISE with Resource Owned Password Credentials, on page 2394
  - About Entra ID and Cisco ISE with TEAP/EAP-TLS, on page 2394
- To download groups from Entra ID.

  For more information about setting up Entra ID, see Configure Microsoft Entra ID for Passive Authentication, on page 2395.

Active authentication

Active authentications occur when a user authenticates through preconfigured managed devices. Captive portal is another name for active authentication. Active authentication generally uses the same user repositories as passive authentication (the exceptions being ISE/ISE-PIC, TS Agent, and the Passive Identity Agent, which are passive only).

To use Microsoft Entra ID as a captive portal requires users to authenticate with Entra ID. We refer to the realm as a Security Assertion Markup Language (SAML) realm because SAML is used to establish a trust relationship between:

- A *service provider* (the Secure Firewall Threat Defense device or devices to which authentication requests are sent).
- An *identity provider* (Microsoft Entra ID).

SAML is an open standard developed by the OASIS standards body; for more information, see the SAML Overview.

For more information, see How to Create a Microsoft Azure AD (SAML) Realm for Active Authentication (Captive Portal), on page 2403.

How to Create a Microsoft Azure AD (SAML) for Passive Authentication

This topic discusses the high-level tasks of creating a Microsoft Azure AD (SAML) (now called *Entra ID*) realm for passive authentication use with the Secure Firewall Management Center.

Procedure

	Command or Action	Purpose
Step 1	Enable the Cisco Secure Dynamic Attributes Connector.	The Cisco Secure Dynamic Attributes Connector is required to use a Microsoft Azure AD (SAML) realm. You can do it first or you can enable it when you create the realm. For more information, see Enable the Cisco Secure Dynamic Attributes Connector, on page 1806.
Step 2	Configure Microsoft Entra ID.	Several configuration tasks are required, including setting up an event hub, giving your application permission to the Microsoft Graph API, and enabling the audit log. See Configure Microsoft Entra ID for Passive Authentication, on page 2395.
Step 3	Configure Cisco ISE.	The way you configure ISE depends on how users authenticate with your system. For more information, see How to Configure ISE for Microsoft Azure AD (SAML), on page 2395.
Step 4	Create a Cisco ISE identity source.	The identity source enables ISE to communicate with the Secure Firewall Management Center.
Step 5	Get the information required to configure your Microsoft Azure AD (SAML) realm.	This information includes client and tenant IDs, client secret, and other information store in Microsoft Entra ID.
Step 6	Configure and verify your realm.	Test the realm's configuration before you start to use it in access control policies. Create a Microsoft Azure AD (SAML) realm as discussed in Create a Microsoft Azure AD (SAML) Realm, on page 2392
Step 7	Create access control policies and rules using your Microsoft Azure AD (SAML) realm.	Unlike other types of realms, you do not need to create an identity policy or associate the identity policy with an access control policy. See Creating a Basic Access Control Policy, on page 1748 and Create and Edit Access Control Rules, on page 1783.

What to do next

See About Entra ID and Cisco ISE with Resource Owned Password Credentials, on page 2394.

About Entra ID and Cisco ISE with Resource Owned Password Credentials

The following figure summarizes an Azure AD (now *called Entra ID*) realm with Cisco ISE and resource owned password credentials (ROPC):

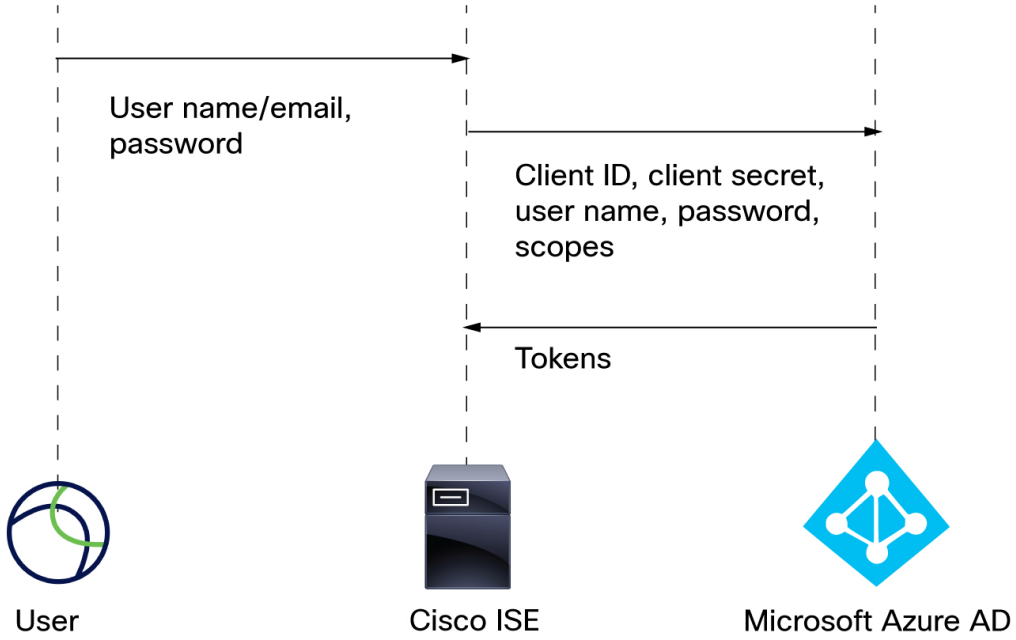

With ROPC,

1. The user logs in with a user name (or email address) and password using a VPN client like Cisco Secure Client.
2. The client ID, client secret, user name, password, and scopes are sent to Entra ID.
3. Tokens are sent from Entra ID to Cisco ISE, which sends user sessions to the Secure Firewall Management Center.

For details about configuring Cisco ISE, see Configure ISE 3.0 REST ID with Azure Active Directory.

Additional resource: Microsoft identity platform and OAuth 2.0 Resource Owner Password Credentials on learn.microsoft.com.

About Entra ID and Cisco ISE with TEAP/EAP-TLS

Tunnel Extensible Authentication Protocol (TEAP) , defined by RFC7170, can be used with ISE and the Secure Firewall Management Center as follows:

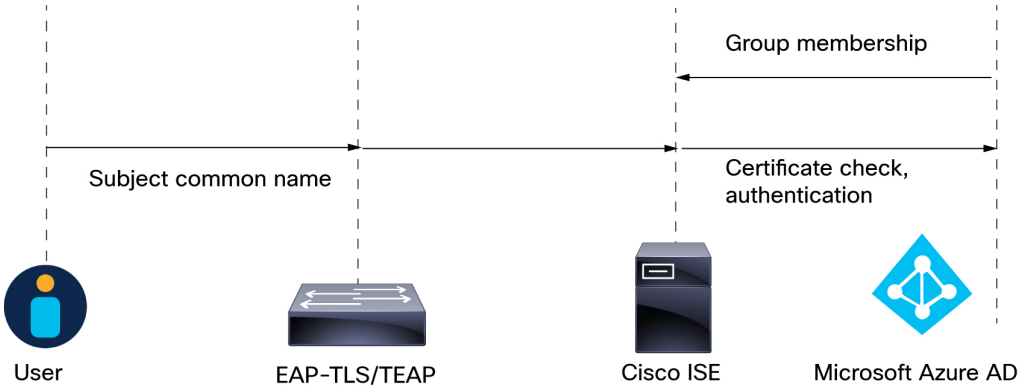

The following is based on Configure Cisco ISE 3.2 EAP-TLS with Microsoft Azure Active Directory:

1. The user's certificate is sent to ISE through EAP-TLS or TEAP with EAP-TLS as the inner method.
2. ISE evaluates the user's certificate (validity period, trusted certificate authority, certificate revocation list, and so on).
3. ISE takes the certificate subject name (CN) and performs a look-up to the Azure Graph API to fetch the user's groups and other attributes. This is referred to by Azure as User Principal name (UPN).
4. ISE authorization policies are evaluated against the user's attributes returned from Azure.

How to Configure ISE for Microsoft Azure AD (SAML)

In a Microsoft Azure AD (SAML) (now called *Entra ID*) realm, it's the responsibility of ISE to send user sessions (login, logout) to the Secure Firewall Management Center. This topic discusses how to set up ISE for use with an Azure AD realm.

Resource owned password credentials authentication

To use ISE with Microsoft Azure AD (SAML) implemented through Representational State Transfer (REST) Identity (ID) service with the help of Resource Owner Password Credentials (ROPC), see Configure ISE 3.0 REST ID with Azure Active Directory.

TEAP/EAP-TLS

To use ISE with authorization policies based on Azure AD group membership and other user attributes with EAP-TLS or TEAP as the authentication protocols, see Configure Cisco ISE 3.2 EAP-TLS with Microsoft Azure Active Directory.

What to do next

Get Required Information For Your Microsoft Azure AD (SAML) Realm, on page 2398

Configure Microsoft Entra ID for Passive Authentication

This topic provides basic information about how to set up a Microsoft Entra ID (formerly caled *Microsoft Azure Active Directory (AD)*) as a realm you can use with the Secure Firewall Management Center. We expect you to already be familiar with Entra ID; if not, consult documentation or a support resource before you get started.

Give your application the Microsoft Graph permission

Grant your Entra ID application the following permissions to Microsoft Graph as discussed in Authorization and the Microsoft Graph Security API on the Microsoft site:

- Reader role
- User.Read.All permission
- Group.Read.All permission

This permission enables the Secure Firewall Management Center to download users and groups from Entra ID the first time.

Required information from this step for setting up the Entra ID realm in the Secure Firewall Management Center:

- Name of the app you registered
- **Application (client) ID**
- **Client secret**
- **Directory (tenant) ID**

Set up an event hub

Set up the event hub as discussed in Quickstart: Create an event hub using Azure portal on the Microsoft site. The Secure Firewall Management Center uses the event hub audit log to download periodic updates to users and groups.

More information: Features and terminology in Azure Event Hubs.

Important You must choose the **Standard** pricing tier or better. If you choose **Basic**, the realm cannot be used.

Required information from this step for setting up the Entra ID realm in the Secure Firewall Management Center:

- Namespace **Name**
- **Connection string—primary key**
- Event Hub **Name**
- Consumer group **Name**

Enable the audit log

Enable the audit log as discussed in Tutorial: Stream Azure Active Directory logs to an Azure event hub on the Microsoft site.

Configure Cisco ISE for Entra ID

To send user session information to the Secure Firewall Management Center, configure Cisco ISE for Entra ID as discussed in Configure ISE 3.0 REST ID with Azure Active Directory.

What to do next

See How to Configure ISE for Microsoft Azure AD (SAML), on page 2395.

Configure Entra ID Basic Settings

Give your application the Microsoft Graph permission

Grant your Entra ID (formerly called *Azure AD*) application the following permissions to Microsoft Graph as discussed in Authorization and the Microsoft Graph Security API on the Microsoft site:

- Reader role
- User.Read.All permission
- Group.Read.All permission

This permission enables the Firewall Management Center to download users and groups from Entra ID the first time.

Required information from this step for setting up the Entra ID realm in the Firewall Management Center:

- Name of the app you registered
- **Application (client) ID**
- **Client secret**
- **Directory (tenant) ID**

Set up an event hub

Set up the event hub as discussed in Quickstart: Create an event hub using Azure portal on the Microsoft site. The Firewall Management Center uses the event hub audit log to download periodic updates to users and groups.

More information: Features and terminology in Azure Event Hubs.

Important You must choose the **Standard** pricing tier or better. If you choose **Basic**, the realm cannot be used.

Required information from this step for setting up the Entra ID realm in the Secure Firewall Management Center:

- Namespace **Name**
- **Connection string—primary key**
- Event Hub **Name**
- Consumer group **Name**

Enable the audit log

Enable the audit log as discussed in Tutorial: Stream Azure Active Directory logs to an Azure event hub on the Microsoft site.

Get Required Information For Your Microsoft Azure AD (SAML) Realm

This task explains how to get the information required to set up a Microsoft Azure AD (SAML) (now called *Entra ID*) realm in the Secure Firewall Management Center. You might have already obtained this information when you set up Microsoft Entra ID as discussed in Configure Microsoft Entra ID for Passive Authentication, on page 2395.

Procedure

Step 1 Log in to https://portal.azure.com/ as a user with at least the Product Designer role.

Step 2 At the top of the page, click **Microsoft Entra ID**.

Step 3 In the left column, click **App Registrations**.

Step 4 If necessary, filter the list of displayed apps to show the one you want to use.

Step 5 Click the name of your app.

Step 6 Click **Copy** ( ) next to the following values on this page and paste those values to a text file.

- **Application (Client) ID**
- **Directory (tenant) ID**

Step 7 Click **Client Credentials**.

Step 8 Unless you already know the client secret *value* (as opposed to the client *secret ID*), you must create a new client secret as follows:

a) Click **New Client Secret**.

b) Enter the required information in the provided fields.

c) Click **Add**.

d) Click **Copy** ( ) next to Value as the following figure shows.

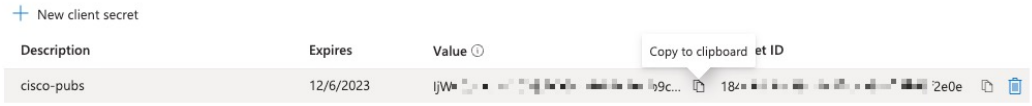

Step 9 From https://portal.azure.com/, click **Event Hubs** > **(name of an event hub)**.

Step 10 In the right pane, click **Copy** ( ) next to the value of **Host name** and paste the value to the clipboard. This is your event hub host name.

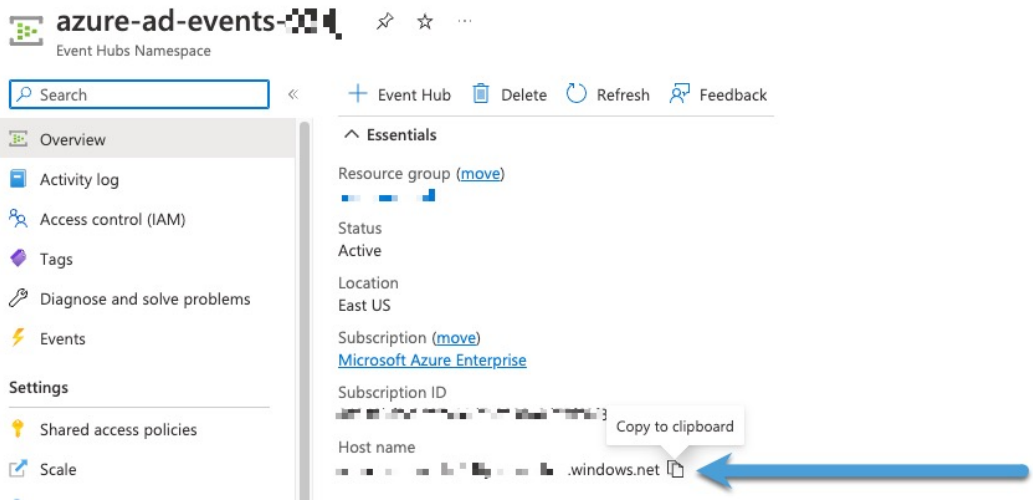

Step 11 Write down or copy to a text file the name of the event hub (same as the **Event Hubs Namespace** at the top of the page).

Step 12 In the left pane, under Settings, click **Shared access policies**.

Step 13 Click the name of a policy.

Step 14 Click **Copy** (⧉) next to **Connection string-primary key**.

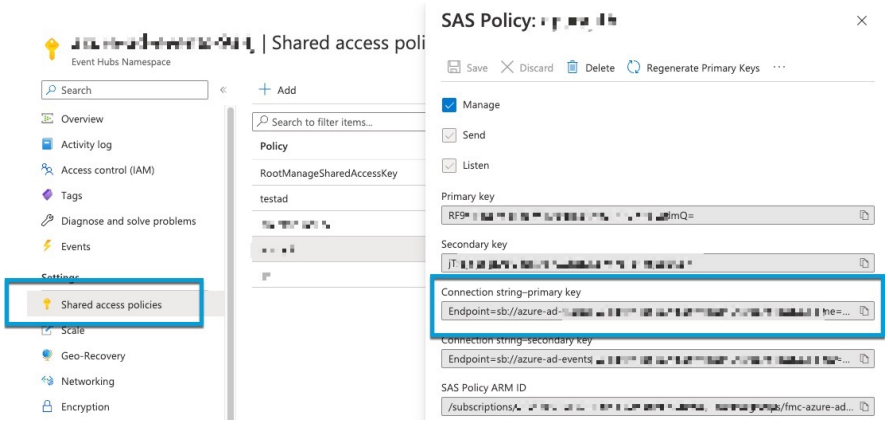

Step 15 Click **Overview** > **Entities** > **Event Hubs** > **(name of an event hub)** > **Entities** > **Consumer Groups**.

Write down the following value or copy it to the clipboard. This is your consumer group name.

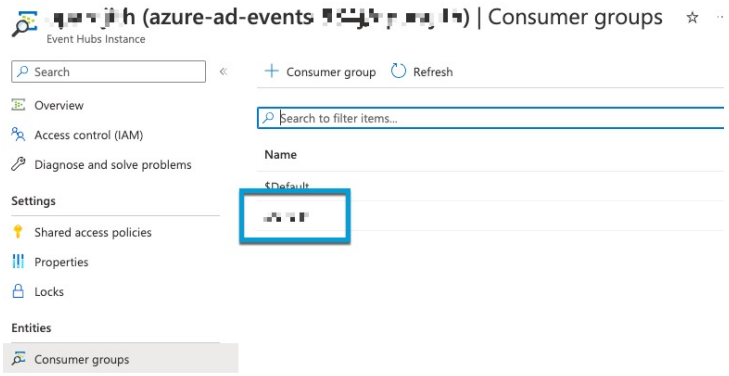

Step 16 In the left pane, click **Overview**.

Step 17 Click **Copy** ( ) next to **Namespace**.

This is your event hubs topic name.

Create a Microsoft Azure AD (SAML) Realm for Passive Authentication

The following topics discuss how to run the multi-step wizard required to create a Microsoft Azure AD (SAML) (now called *Entra ID*) realm for passive authentication.

You can use a Microsoft Azure Active Directory (AD) realm with Cisco ISE to authenticate users and get user sessions for user control. We get groups from Entra ID and logged-in user session data from Cisco ISE.

You have the following options:

- Resource owned password credentials (ROPC): Enables users to log in with a client like Cisco Secure Client using a user name and password. ISE sends user sessions to the Secure Firewall Management Center. For more information, see About Entra ID and Cisco ISE with Resource Owned Password Credentials, on page 2394.

  Additional resource: Microsoft identity platform and OAuth 2.0 Resource Owner Password Credentials on learn.microsoft.com.

- Extensible Authentication Protocol (EAP) Chaining with Tunnel-based Extensible Authentication Protocol (TEAP) and Transport Layer Security (TLS), abbreviated EAP/TEAP-TLS: TEAP is a tunnel-based EAP method that establishes a secure tunnel and executes other EAP methods under the protection of that secured tunnel. ISE is used to validate user credentials and to send user sessions to the Secure Firewall Management Center. For more information, see About Entra ID and Cisco ISE with TEAP/EAP-TLS, on page 2394.

To configure the realm, complete all tasks in the following order:

1. Configure Entra ID Basic Settings, on page 2397.

2. Get required information for your realm as discussed in Get Required Information For Your Microsoft Azure AD (SAML) Realm, on page 2398.
3. Microsoft Azure AD (SAML) Realm: SAML Details, on page 2401.

Microsoft Azure AD (SAML) Realm: SAML Details

This task discusses the first step in a multi-step wizard that creates a Microsoft Azure AD (SAML) realm. You must complete all steps in the wizard to set up your realm. The steps are different depending on whether you create the realm for active or passive authentication.

Before you begin

Complete all of the following tasks before you create your realm:

- (Passive authentication with Cisco ISE only.) If you're using Cisco ISE as the repository for users and groups, set up ISE:
  - About Entra ID and Cisco ISE with Resource Owned Password Credentials, on page 2394
  - About Entra ID and Cisco ISE with TEAP/EAP-TLS, on page 2394
- To use Entra ID as the repository for users and groups, see Configure Microsoft Entra ID for Passive Authentication, on page 2395.
- Get required information for your realm as discussed in Get Required Information For Your Microsoft Azure AD (SAML) Realm, on page 2398.

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Integration** > **Other Integrations** > **Realms** > **Realms**.

Step 3 Click **Add Realm** > **SAML - Azure AD**.

Step 4 Enter the following information.

Item	Description
Name	Unique name to identify the realm.
Description	(Optional.) Description of the realm.
Identity Provider	Always displays **Azure AD**.
Configuration Type	Click one of the following: • **Passive Authentication with ISE** for passive authentication. • **Passive authentication or captive portal with Azure AD** for to use Entra ID as a user store for either passive authentication or active authentication (that is, captive portal).

Step 5 Click **Next**.

What to do next

One of the following:

- Passive authentication: Microsoft Azure AD (SAML) Realm: Azure AD Details, on page 2402.
- Active authentication: Microsoft Azure AD (SAML) Realm: SAML Service Provider (SP) Metadata, on page 2414.

Microsoft Azure AD (SAML) Realm: Azure AD Details

This task discussed one page in a multi-page wizard to enable you to create a Microsoft Azure AD (SAML) realm.

Procedure

Step 1 Continue from the preceding step in the wizard.

Step 2 Enter the following information.

Item	Description
Name	Enter a unique name to identify this realm.
Client Secret **Tenant ID** **Event Hubs Host Name** **Event Hub Name** **Event Hub Connection String**	Enter the information you found as discussed in:. • Passive authentication: How to Create a Microsoft Azure AD (SAML) for Passive Authentication, on page 2393 • Active authentication: Get Required Information For Your Microsoft Azure AD (SAML) Realm (Active Authentication Only), on page 2409
(Optional.) **User Groups**	Slide to **Slider enabled** ( ) to specify groups to include or exclude from policy.
(Optional.) **Excluded User Groups**	If you enter one or more group names in this field, all groups and users they contain *except* these are downloaded and available for user awareness and user control. Enter one group name per line followed by a line break. Group names are case-sensitive.
(Optional.) **Included User Groups**	If you enter one or more group names in this field, only those groups and users they contain are downloaded and user data is available for user awareness and user control. Enter one group name per line followed by a line break. Group names are case-sensitive.

Step 3 Click **Test**.
Make sure the test connection succeeds before you continue to the next step.

Step 4 Click **Next**.

Microsoft Azure AD (SAML) Realm: User Session Timeout

This task discussed one page in a multi-page wizard to enable you to create a Microsoft Azure AD (SAML) realm.

This option sets the number of seconds before an inactive session is terminated by the system.

Procedure

Step 1 Continue from the preceding step in the wizard.

Step 2 Enter the following information.

Item	Description
ISE users	Default is 1440 minutes (24 hours).
Captive portal users	Default is 1440 minutes (24 hours).

After the timeout is exceeded, the user's session ends; if the user continues to access the network without logging in again, the user is seen by the Firewall Management Center as Unknown (except for **Failed Captive Portal Users**).

Step 3 Click **Save**.

How to Create a Microsoft Azure AD (SAML) Realm for Active Authentication (Captive Portal)

This topic discusses the high-level tasks of creating a Microsoft Azure Active Directory (AD) realm (now called *Entra ID*) for use with the Secure Firewall Management Center.

Before you begin

If you enabled Change Management, you must open or edit, assign, and approve a ticket for each of the following objects before you can create the realm:

- Base URL
- Service provider certificate enrollment (PKCS12 format)
- Identity provider certificate enrollment (manual format)
- The realm itself (create and assign the ticket until realm creation is complete, then approve it)

For more information, see Opening a Ticket for Configuration Changes, on page 241 and Policies and Objects that Support Change Management, on page 235.

Procedure

	Command or Action	Purpose
Step 1	Create a fully-qualified host name (FQDN) using your DNS server and upload the Firewall Threat Defense's internal certificate to the Secure Firewall Management Center. You can consult a resource such as this one if you've never done it before. Specify the IP address of a routed interface on one of the devices managed by your Secure Firewall Management Center.	Consult a DNS server reference.
Step 2	Enable the Cisco Secure Dynamic Attributes Connector.	The Cisco Secure Dynamic Attributes Connector is required to use a Microsoft Azure AD realm. You can do it first or you can enable it when you create the realm. For more information, see Enable the Cisco Secure Dynamic Attributes Connector, on page 1806.
Step 3	Create a network object with an associated internal certificate.	See Creating Network Objects, on page 1403.
Step 4	Get a signed certificate and upload it to the Secure Firewall Threat Defense to which Entra ID authentication requests will be sent.	The certificate should be signed by a trusted Certificate Authority (CA) and delivered to you in **.p12** format (also referred to as PKCS#12; see also this article on ssl.com). For background, see the section on public key infrastructure in Cisco Secure Firewall Management Center Device Configuration Guide or stackoverflow.com. To upload the signed certificate, see Installing a Certificate Using a PKCS12 File, on page 1492.
Step 5	Configure Microsoft Entra ID basic settings.	Several configuration tasks are required, including setting up an event hub, giving your application permission to the Microsoft Graph API, and enabling the audit log. See Configure Entra ID Basic Settings, on page 2397.
Step 6	Create a single sign-on (SSO) app in Entra ID.	The SSO app enables users that request access to a protected network resource to authenticate with Entra ID. The SSO app has both the federation XML that you can use to simplify realm creation as well as the identity provider

	Command or Action	Purpose
		certificate the Secure Firewall Threat Defense requires to security authenticate with Entra ID. See Configure a Single Sign-On (SSO) App in Entra ID, on page 2407.
Step 7	Get the information required to configure your Microsoft Azure AD (SAML) realm.	This information includes client and tenant IDs, client secret, and other information store in Microsoft Entra ID. See Get Required Information For Your Microsoft Azure AD (SAML) Realm, on page 2398.
Step 8	Configure a decryption policy with a **Decrypt - Resign** rule for the Azure Authentication Service so users can access web pages using the HTTPS protocol.	The Microsoft Azure AD (SAML) realm can authenticate users only if the HTTPS traffic is decrypted before the traffic is sent to the realm. The Microsoft Azure AD (SAML) realm itself is seen by the system as the **Azure Authentication Service** application. Create a Decryption Rule with Decrypt - Resign Action, on page 2408.
Step 9	Create identity policy with an active authentication rule.	The identity policy enables selected users in your realm access resources after authenticating with the Microsoft Azure AD (SAML) realm. For more information, see Create an Identity Policy, on page 2555.
Step 10	Create access control policies and rules using your Microsoft Azure AD realm.	Unlike other types of realms, you do not need to create an identity policy or associate the identity policy with an access control policy. See Creating a Basic Access Control Policy, on page 1748 and Create and Edit Access Control Rules, on page 1783.
Step 11	Associate the identity and decryption policies with the access control policy from step 3.	This final step enables the system to authenticate users with the Microsoft Azure AD (SAML) realm. For more information, see Associating Other Policies with Access Control, on page 1762.

What to do next

See Configure Entra ID Basic Settings, on page 2397.

Configure Entra ID Basic Settings

Give your application the Microsoft Graph permission

Grant your Entra ID (formerly called *Azure AD*) application the following permissions to Microsoft Graph as discussed in Authorization and the Microsoft Graph Security API on the Microsoft site:

- Reader role
- User.Read.All permission
- Group.Read.All permission

This permission enables the Firewall Management Center to download users and groups from Entra ID the first time.

Required information from this step for setting up the Entra ID realm in the Firewall Management Center:

- Name of the app you registered
- **Application (client) ID**
- **Client secret**
- **Directory (tenant) ID**

Set up an event hub

Set up the event hub as discussed in Quickstart: Create an event hub using Azure portal on the Microsoft site. The Firewall Management Center uses the event hub audit log to download periodic updates to users and groups.

More information: Features and terminology in Azure Event Hubs.

Important You must choose the **Standard** pricing tier or better. If you choose **Basic**, the realm cannot be used.

Required information from this step for setting up the Entra ID realm in the Secure Firewall Management Center:

- Namespace **Name**
- **Connection string—primary key**
- Event Hub **Name**
- Consumer group **Name**

Enable the audit log

Enable the audit log as discussed in Tutorial: Stream Azure Active Directory logs to an Azure event hub on the Microsoft site.

Configure a Single Sign-On (SSO) App in Entra ID

This topic discusses how to create an app in Microsoft Entra ID to handle single sign-on (SSO) from Entra ID when a network user attempts to access a protected network.

Create the app

In the Microsoft Entra ID portal, click **Enterprise Applications** on the home page and follow the instructions in Configure Microsoft Entra SSO on learn.microsoft.com.

The following figure shows part of the SSO app configuration. You must provide some of the information on this page when you configure the Microsoft Azure AD (SAML) realm. For more information, see Get Required Information For Your Microsoft Azure AD (SAML) Realm (Active Authentication Only), on page 2409.

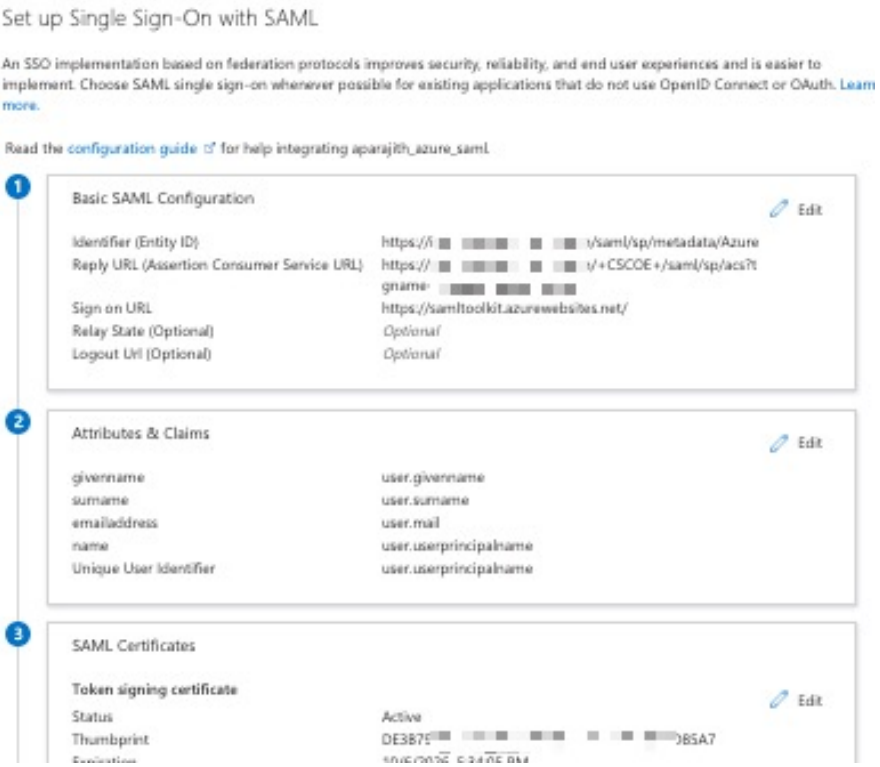

(Optional.) Upload the service provider metadata

If you already configured the Microsoft Azure AD (SAML) realm, click Upload metadata file at the top of the page to quickly provide configuration values for the SSO app.

The following figure shows an example.

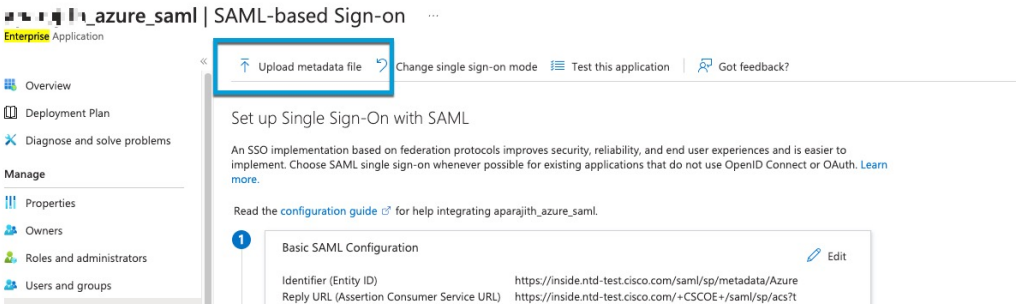

Add users and groups to the SSO app

Add users and groups to your app as discussed in Add a user account to an enterprise application on learn.microsoft.com

What to do next

See Get Required Information For Your Microsoft Azure AD (SAML) Realm (Active Authentication Only), on page 2409.

Create a Decryption Rule with Decrypt - Resign Action

This part of the procedure discusses how to create a decryption policy to decrypt and resign traffic before the traffic reaches the SAML realm. The realm can authenticate traffic only after it has been decrypted.

Procedure

Step 1 If you haven't done so already, log in to the Secure Firewall Management Center.

Step 2 If you haven't done so already, create an internal certificate authority object to decrypt TLS/SSL traffic as discussed in PKI, on page 1404.

Step 3 Click **Policies** > **Access Control heading** > **Decryption**.

Step 4 Click **New Policy**.

Step 5 Enter a **Name** and choose a **Default Action** for the policy. Default actions are discussed in Decryption Policy Default Actions, on page 2295.

Step 6 Click **Save**.

Step 7 Click **Add Rule**.

Step 8 Enter a **Name** for the rule.

Step 9 From the **Action** list, choose **Decrypt - Resign**.

Step 10 From the **with** list, choose your service provider certificate object.

Step 11 Click the **Applications** tab page.

Step 12 In the Available Applications section, enter **Azure Authentication Service** in the search field.

Step 13 Click **Azure Authentication** and click **Add to Rule**.
The following figure shows an example.

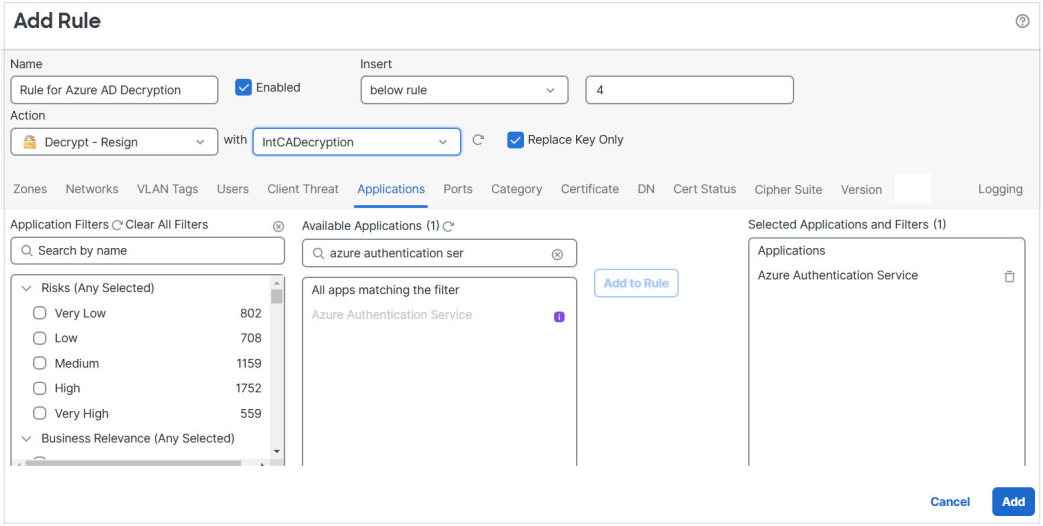

Step 14 (Optional.) Set other options as discussed in Decryption Rule Conditions, on page 2316.

Step 15 Click **Add**.

Step 16 At the top of the page, click **Save**.

What to do next

Get Required Information For Your Microsoft Azure AD (SAML) Realm (Active Authentication Only)

This task explains how to get the information required to set up a Microsoft Azure AD (SAML) realm (now called *Entra ID*) in the Firewall Management Center.

Procedure

Step 1 Log in to https://portal.azure.com/ as a user with at least the Product Designer role.

Step 2 At the top of the page, click **Microsoft Entra ID**.

Step 3 In the left column, click **App Registrations**.

Step 4 If necessary, filter the list of displayed apps to show the one you want to use.

Step 5 Click the name of your app.

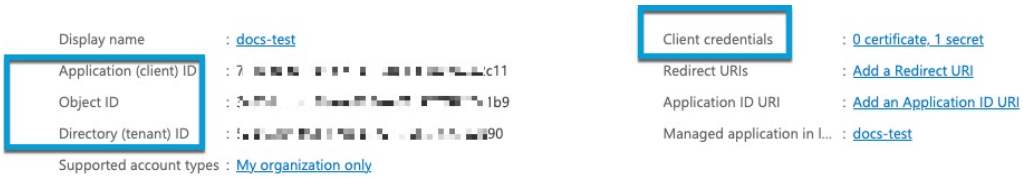

Step 6 Click **Copy** ( ) next to the following values on this page and paste those values to a text file.

- **Application (Client) ID**

- **Directory (tenant) ID**

Step 7 Click **Client Credentials**.

Step 8 Unless you already know the client secret *value* (as opposed to the client *secret ID*), you must create a new client secret as follows:

a) Click **New Client Secret**.

b) Enter the required information in the provided fields.

c) Click **Add**.

d) Click **Copy** (⧉) next to Value as the following figure shows.

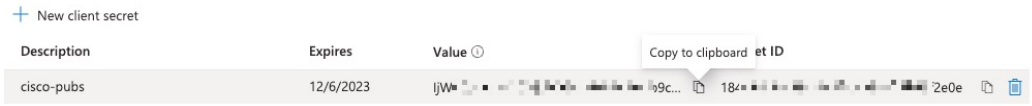

Step 9 From https://portal.azure.com/, click **Event Hubs** > **(name of an event hub)**.

Step 10 In the right pane, click **Copy** (⧉) next to the value of **Host name** and paste the value to the clipboard. This is your event hub host name.

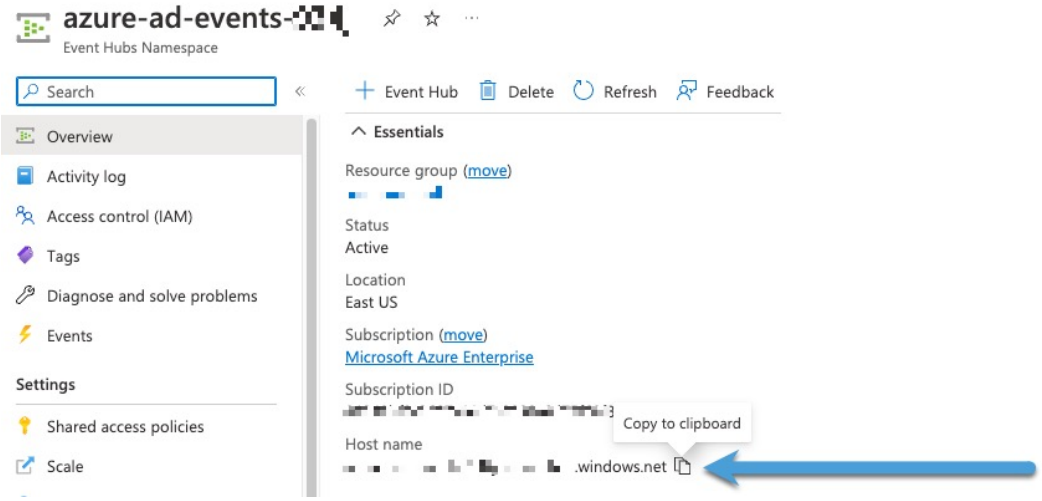

Step 11 Write down or copy to a text file the name of the event hub (same as the **Event Hubs Namespace** at the top of the page).

Step 12 In the left pane, under Settings, click **Shared access policies**.

Step 13 Click the name of a policy.

Step 14 Click **Copy** (⧉) next to **Connection string-primary key**.

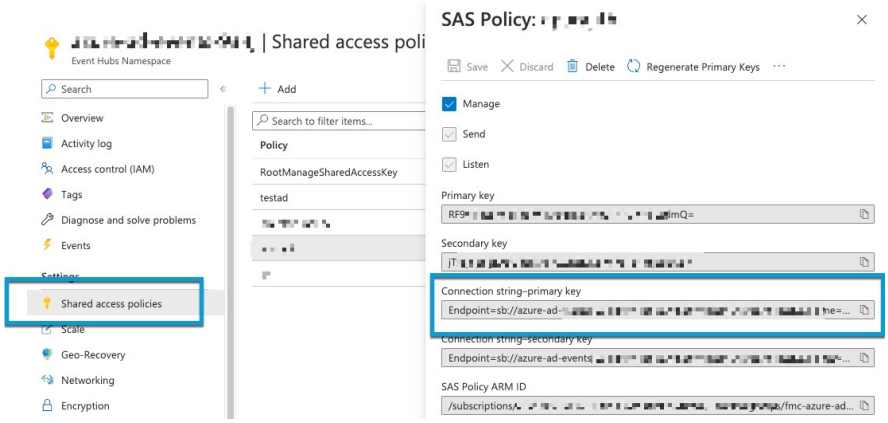

Step 15 Click **Overview** > **Entities** > **Event Hubs** > **(name of an event hub)** > **Entities** > **Consumer Groups**.

Write down the following value or copy it to the clipboard. This is your consumer group name.

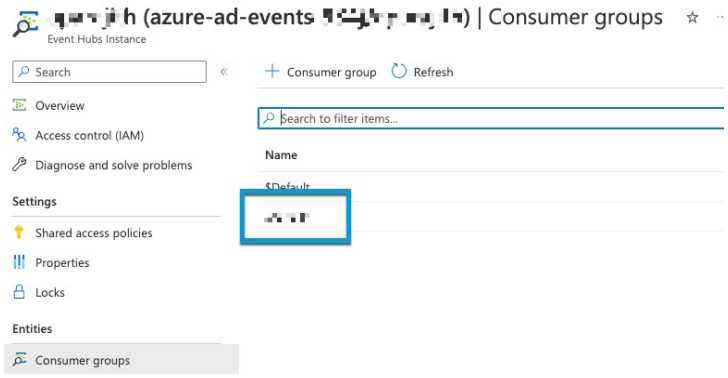

Step 16 In the left pane, click **Overview**.

Step 17 Click **Copy** ( ) next to **Namespace**.

This is your event hubs topic name.

Step 18 Return to the home page and log in if necessary: https://portal.azure.com/#home.

Step 19 Click **Microsoft Entra ID**.

Step 20 In the left pane, click **Enterprise Applications.**

Step 21 If necessary, filter the list of applications to locate yours.

Step 22 Click the name of your enterprise application.

Step 23 Click **Get Started** under Set up single sign on.

Step 24 On your SSO app page, copy the value of Identifier (Entity ID) to the clipboard.
The following figure shows an example.

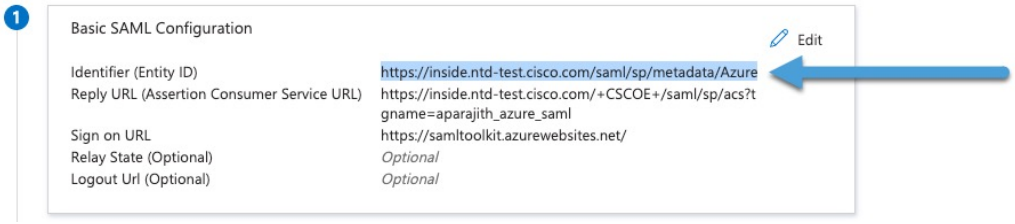

Step 25 On your SSO app page, click the **Download** link next to Federation Metadata XML, as the following figure shows.
The following figure shows an example.

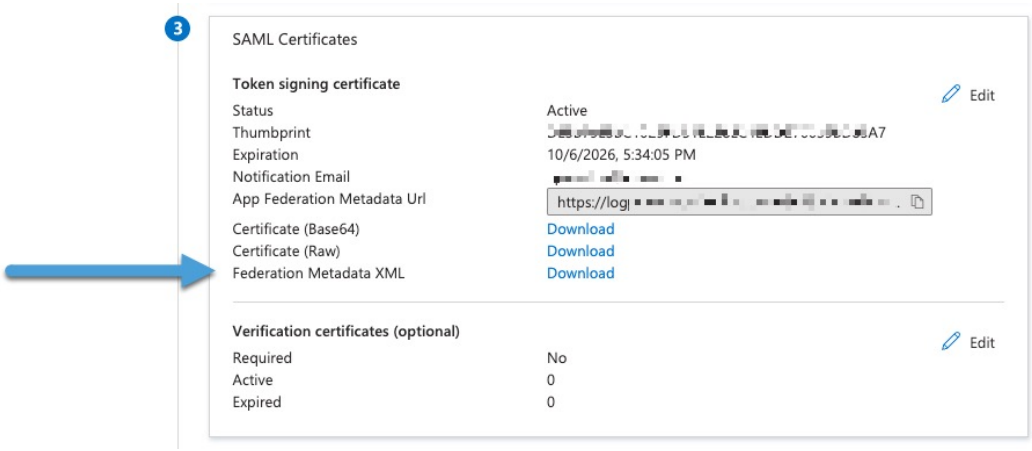

Step 26 If you've already set up your SSO app, *you can stop here*. The Federation Metadata XML contains all the information required to configure the identity provider in the Secure Firewall Management Center.

Step 27 (Optional if you already downloaded the federation XML.) Click **Copy** ( ) next to both of the following values and save them to a text file.
The following figure shows an example.

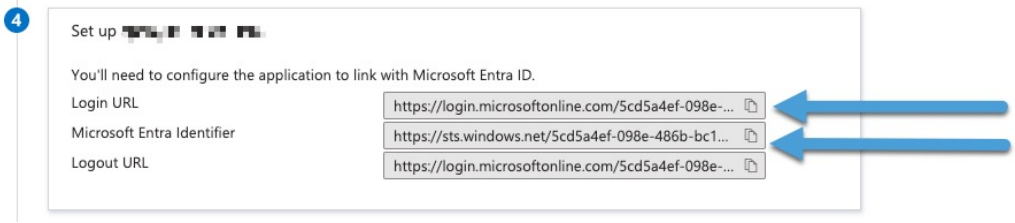

What to do next

See Create a Decryption Rule with Decrypt - Resign Action, on page 2408.

Create a Microsoft Azure AD (SAML) Realm for Active Authentication (Captive Portal)

The following topics discuss how to run the multi-step wizard required to create a Microsoft Azure AD (SAML) realm (now called *Entra ID*) for active authentication.

In active authentication (also referred to as *captive portal*), Microsoft Entra ID is the user store. When a user tries to access a protected resource as defined in access control rules, the user must authenticate with Microsoft Entra ID first.

To configure the realm, complete all tasks in the following order:

1. Configure Entra ID Basic Settings, on page 2397.
2. Configure a Single Sign-On (SSO) App in Entra ID, on page 2407.
3. Get Required Information For Your Microsoft Azure AD (SAML) Realm (Active Authentication Only), on page 2409

Microsoft Azure AD (SAML) Realm: SAML Details

This task discusses the first step in a multi-step wizard that creates a Microsoft Azure AD (SAML) realm. You must complete all steps in the wizard to set up your realm. The steps are different depending on whether you create the realm for active or passive authentication.

Before you begin

Complete all of the following tasks before you create your realm:

- (Passive authentication with Cisco ISE only.) If you're using Cisco ISE as the repository for users and groups, set up ISE:
  - About Entra ID and Cisco ISE with Resource Owned Password Credentials, on page 2394
  - About Entra ID and Cisco ISE with TEAP/EAP-TLS, on page 2394
- To use Entra ID as the repository for users and groups, see Configure Microsoft Entra ID for Passive Authentication, on page 2395.
- Get required information for your realm as discussed in Get Required Information For Your Microsoft Azure AD (SAML) Realm, on page 2398.

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Integration** > **Other Integrations** > **Realms** > **Realms**.

Step 3 Click **Add Realm** > **SAML - Azure AD**.

Step 4 Enter the following information.

Item	Description
Name	Unique name to identify the realm.
Description	(Optional.) Description of the realm.
Identity Provider	Always displays **Azure AD**.

Item	Description
Configuration Type	Click one of the following: • **Passive Authentication with ISE** for passive authentication. • **Passive authentication or captive portal with Azure AD** for to use Entra ID as a user store for either passive authentication or active authentication (that is, captive portal).

Step 5 Click **Next**.

What to do next

One of the following:

- Passive authentication: Microsoft Azure AD (SAML) Realm: Azure AD Details, on page 2402.
- Active authentication: Microsoft Azure AD (SAML) Realm: SAML Service Provider (SP) Metadata, on page 2414.

Microsoft Azure AD (SAML) Realm: SAML Service Provider (SP) Metadata

This task discussed one page in a multi-page wizard to enable you to create a Microsoft Azure AD (SAML) realm.

Before you begin

Complete the tasks discussed in Microsoft Azure AD (SAML) Realm: SAML Details, on page 2401

Procedure

Step 1 Continue from Microsoft Azure AD (SAML) Realm: SAML Details, on page 2401.

Step 2 Enter the following information.

Item	Description
Base URL	From the list, click the network object you previously created. Network users are directed to this URL when they try to access protected network resources. You can also click **Add** (+) to create an object now.
Entity ID	Your SSO app's entity ID.
Assertion Consumer Services (ACS) URL	Automatically generated from the preceding values.

Item	Description
Service Provider Certificate	From the list, click the certificate to use to decrypt requests to the Secure Firewall Threat Defense. You can also click **Add** (+) to create an object now.
Download Service Provider Metadata	(Optional.) Download the metadata associated with the service provider (that is, managed device) to simplify configuring your Microsoft Entra ID SSO app.

Step 3 Click **Next**.

What to do next

Microsoft Azure AD (SAML) Realm: SAML Identity Provider (IdP) Metadata, on page 2415.

Microsoft Azure AD (SAML) Realm: SAML Identity Provider (IdP) Metadata

This task discussed one page in a multi-page wizard to enable you to create a Microsoft Azure AD (SAML) realm.

Before you begin

Complete the tasks discussed in Microsoft Azure AD (SAML) Realm: SAML Service Provider (SP) Metadata, on page 2414.

Procedure

Step 1 Continue from Microsoft Azure AD (SAML) Realm: SAML Service Provider (SP) Metadata, on page 2414.

Step 2 If you previously downloaded the Entra ID SSO app's federation XML, click **Upload XML** and upload it now. You can then skip the next step.

Step 3 Enter the following information.

Item	Description
Entity ID	Enter your identity provider's entity ID.
Single Sign on (SSO) URL	Enter the app's SSO URL.
IdP Certificate	From the list, click the certificate to use to authenticate with Microsoft Entra ID. You can also click **Add** (+) to create an object now.

Step 4 Click **Next**.

What to do next

Microsoft Azure AD (SAML) Realm: SAML Details, on page 2401.

Microsoft Azure AD (SAML) Realm: Azure AD Details

This task discussed one page in a multi-page wizard to enable you to create a Microsoft Azure AD (SAML) realm.

Procedure

Step 1 Continue from the preceding step in the wizard.

Step 2 Enter the following information.

Item	Description
Name	Enter a unique name to identify this realm.
Client Secret	Enter the information you found as discussed in:. • Passive authentication: How to Create a Microsoft Azure AD (SAML) for Passive Authentication, on page 2393 • Active authentication: Get Required Information For Your Microsoft Azure AD (SAML) Realm (Active Authentication Only), on page 2409
Tenant ID	
Event Hubs Host Name	
Event Hub Name	
Event Hub Connection String	
(Optional.) **User Groups**	Slide to **Slider enabled** ( ) to specify groups to include or exclude from policy.
(Optional.) **Excluded User Groups**	If you enter one or more group names in this field, all groups and users they contain *except* these are downloaded and available for user awareness and user control. Enter one group name per line followed by a line break. Group names are case-sensitive.
(Optional.) **Included User Groups**	If you enter one or more group names in this field, only those groups and users they contain are downloaded and user data is available for user awareness and user control. Enter one group name per line followed by a line break. Group names are case-sensitive.

Step 3 Click **Test**.
Make sure the test connection succeeds before you continue to the next step.

Step 4 Click **Next**.

Microsoft Azure AD (SAML) Realm: User Session Timeout

This task discussed one page in a multi-page wizard to enable you to create a Microsoft Azure AD (SAML) realm.

This option sets the number of seconds before an inactive session is terminated by the system.

Procedure

Step 1 Continue from the preceding step in the wizard.

Step 2 Enter the following information.

Item	Description
ISE users	Default is 1440 minutes (24 hours).
Captive portal users	Default is 1440 minutes (24 hours).

After the timeout is exceeded, the user's session ends; if the user continues to access the network without logging in again, the user is seen by the Firewall Management Center as Unknown (except for **Failed Captive Portal Users**).

Step 3 Click **Save**.

Create a Microsoft Azure AD (SAML) Realm for Passive Authentication

You can use a Microsoft Azure Active Directory (AD) realm with ISE to authenticate users and get user sessions for user control. We get groups from Azure AD and logged-in user session data from ISE.

You have the following options:

- Resource owned password credentials (ROPC): Enables users to log in with a client like AnyConnect using a user name and password. ISE sends user sessions to the Secure Firewall Management Center. For more information, see About Entra ID and Cisco ISE with Resource Owned Password Credentials, on page 2394.

  Additional resource: Microsoft identity platform and OAuth 2.0 Resource Owner Password Credentials on learn.microsoft.com.

- Extensible Authentication Protocol (EAP) Chaining with Tunnel-based Extensible Authentication Protocol (TEAP) and Transport Layer Security (TLS), abbreviated EAP/TEAP-TLS: TEAP is a tunnel-based EAP method that establishes a secure tunnel and executes other EAP methods under the protection of that secured tunnel. ISE is used to validate user credentials and to send user sessions to the Secure Firewall Management Center. For more information, see About Entra ID and Cisco ISE with TEAP/EAP-TLS, on page 2394.

This type of authentication is *passive* authentication as opposed to active authentication used in Create a Microsoft Azure AD (SAML) Realm, on page 2392.

Note Before deploying policies related to a Microsoft Azure AD realm, see User Limits for Microsoft Azure Active Directory Realms, on page 2386.

About Entra ID and Cisco ISE with Resource Owned Password Credentials

The following figure summarizes an Azure AD (now *called Entra ID*) realm with Cisco ISE and resource owned password credentials (ROPC):

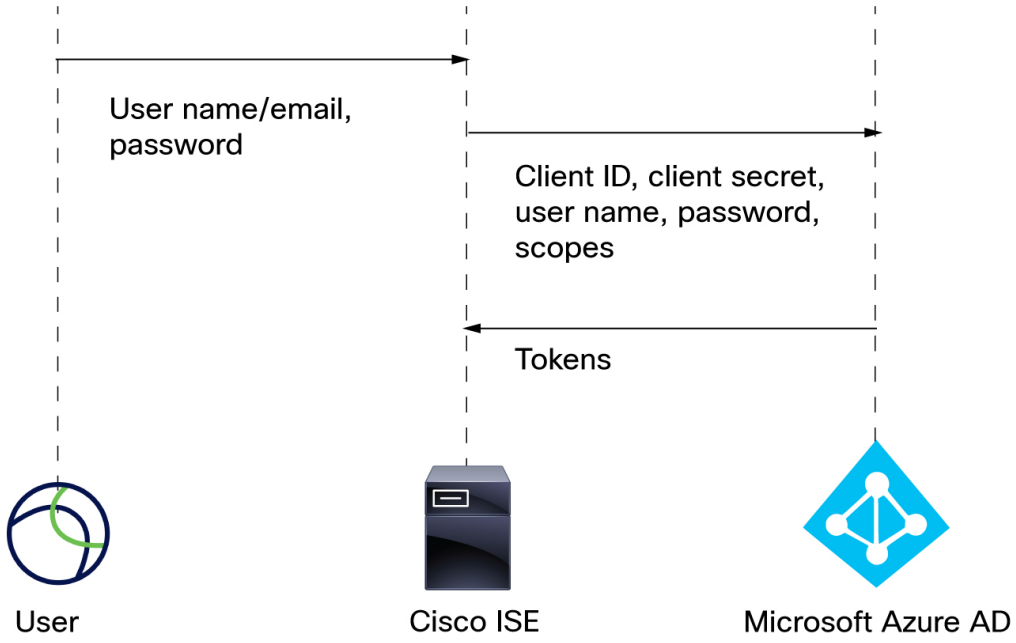

With ROPC,

1. The user logs in with a user name (or email address) and password using a VPN client like Cisco Secure Client.
2. The client ID, client secret, user name, password, and scopes are sent to Entra ID.
3. Tokens are sent from Entra ID to Cisco ISE, which sends user sessions to the Secure Firewall Management Center.

For details about configuring Cisco ISE, see Configure ISE 3.0 REST ID with Azure Active Directory.

Additional resource: Microsoft identity platform and OAuth 2.0 Resource Owner Password Credentials on learn.microsoft.com.

About Entra ID and Cisco ISE with TEAP/EAP-TLS

Tunnel Extensible Authentication Protocol (TEAP) , defined by RFC7170, can be used with ISE and the Secure Firewall Management Center as follows:

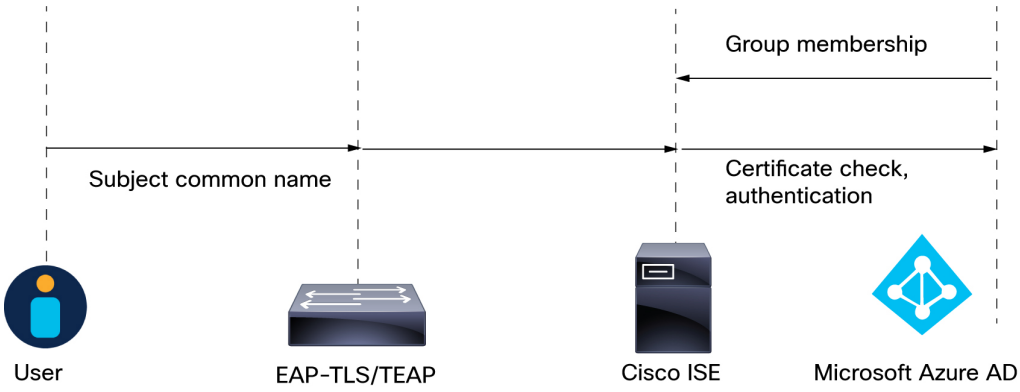

The following is based on Configure Cisco ISE 3.2 EAP-TLS with Microsoft Azure Active Directory:

1. The user's certificate is sent to ISE through EAP-TLS or TEAP with EAP-TLS as the inner method.
2. ISE evaluates the user's certificate (validity period, trusted certificate authority, certificate revocation list, and so on).
3. ISE takes the certificate subject name (CN) and performs a look-up to the Azure Graph API to fetch the user's groups and other attributes. This is referred to by Azure as User Principal name (UPN).
4. ISE authorization policies are evaluated against the user's attributes returned from Azure.

How to Create a Microsoft Azure AD (SAML) for Passive Authentication

This topic discusses the high-level tasks of creating a Microsoft Azure AD (SAML) (now called *Entra ID*) realm for passive authentication use with the Secure Firewall Management Center.

Procedure

	Command or Action	Purpose
Step 1	Enable the Cisco Secure Dynamic Attributes Connector.	The Cisco Secure Dynamic Attributes Connector is required to use a Microsoft Azure AD (SAML) realm. You can do it first or you can enable it when you create the realm. For more information, see Enable the Cisco Secure Dynamic Attributes Connector, on page 1806.
Step 2	Configure Microsoft Entra ID.	Several configuration tasks are required, including setting up an event hub, giving your application permission to the Microsoft Graph API, and enabling the audit log. See Configure Microsoft Entra ID for Passive Authentication, on page 2395.

	Command or Action	Purpose
Step 3	Configure Cisco ISE.	The way you configure ISE depends on how users authenticate with your system. For more information, see How to Configure ISE for Microsoft Azure AD (SAML), on page 2395.
Step 4	Create a Cisco ISE identity source.	The identity source enables ISE to communicate with the Secure Firewall Management Center.
Step 5	Get the information required to configure your Microsoft Azure AD (SAML) realm.	This information includes client and tenant IDs, client secret, and other information store in Microsoft Entra ID.
Step 6	Configure and verify your realm.	Test the realm's configuration before you start to use it in access control policies. Create a Microsoft Azure AD (SAML) realm as discussed in Create a Microsoft Azure AD (SAML) Realm, on page 2392
Step 7	Create access control policies and rules using your Microsoft Azure AD (SAML) realm.	Unlike other types of realms, you do not need to create an identity policy or associate the identity policy with an access control policy. See Creating a Basic Access Control Policy, on page 1748 and Create and Edit Access Control Rules, on page 1783.

What to do next

See About Entra ID and Cisco ISE with Resource Owned Password Credentials, on page 2394.

Configure Microsoft Entra ID for Passive Authentication

This topic provides basic information about how to set up a Microsoft Entra ID (formerly caled *Microsoft Azure Active Directory (AD)*) as a realm you can use with the Secure Firewall Management Center. We expect you to already be familiar with Entra ID; if not, consult documentation or a support resource before you get started.

Give your application the Microsoft Graph permission

Grant your Entra ID application the following permissions to Microsoft Graph as discussed in Authorization and the Microsoft Graph Security API on the Microsoft site:

- Reader role
- User.Read.All permission
- Group.Read.All permission

This permission enables the Secure Firewall Management Center to download users and groups from Entra ID the first time.

Required information from this step for setting up the Entra ID realm in the Secure Firewall Management Center:

- Name of the app you registered
- **Application (client) ID**
- **Client secret**
- **Directory (tenant) ID**

Set up an event hub

Set up the event hub as discussed in Quickstart: Create an event hub using Azure portal on the Microsoft site. The Secure Firewall Management Center uses the event hub audit log to download periodic updates to users and groups.

More information: Features and terminology in Azure Event Hubs.

Important You must choose the **Standard** pricing tier or better. If you choose **Basic**, the realm cannot be used.

Required information from this step for setting up the Entra ID realm in the Secure Firewall Management Center:

- Namespace **Name**
- **Connection string—primary key**
- Event Hub **Name**
- Consumer group **Name**

Enable the audit log

Enable the audit log as discussed in Tutorial: Stream Azure Active Directory logs to an Azure event hub on the Microsoft site.

Configure Cisco ISE for Entra ID

To send user session information to the Secure Firewall Management Center, configure Cisco ISE for Entra ID as discussed in Configure ISE 3.0 REST ID with Azure Active Directory.

What to do next

See How to Configure ISE for Microsoft Azure AD (SAML), on page 2395.

How to Configure ISE for Microsoft Azure AD (SAML)

In a Microsoft Azure AD (SAML) (now called *Entra ID*) realm, it's the responsibility of ISE to send user sessions (login, logout) to the Secure Firewall Management Center. This topic discusses how to set up ISE for use with an Azure AD realm.

Resource owned password credentials authentication

To use ISE with Microsoft Azure AD (SAML) implemented through Representational State Transfer (REST) Identity (ID) service with the help of Resource Owner Password Credentials (ROPC), see Configure ISE 3.0 REST ID with Azure Active Directory.

TEAP/EAP-TLS

To use ISE with authorization policies based on Azure AD group membership and other user attributes with EAP-TLS or TEAP as the authentication protocols, see Configure Cisco ISE 3.2 EAP-TLS with Microsoft Azure Active Directory.

What to do next

Get Required Information For Your Microsoft Azure AD (SAML) Realm, on page 2398

Get Required Information For Your Microsoft Azure AD (SAML) Realm

This task explains how to get the information required to set up a Microsoft Azure AD (SAML) (now called *Entra ID*) realm in the Secure Firewall Management Center. You might have already obtained this information when you set up Microsoft Entra ID as discussed in Configure Microsoft Entra ID for Passive Authentication, on page 2395.

Procedure

Step 1 Log in to https://portal.azure.com/ as a user with at least the Product Designer role.

Step 2 At the top of the page, click **Microsoft Entra ID**.

Step 3 In the left column, click **App Registrations**.

Step 4 If necessary, filter the list of displayed apps to show the one you want to use.

Step 5 Click the name of your app.

Step 6 Click **Copy** ( ) next to the following values on this page and paste those values to a text file.

- **Application (Client) ID**
- **Directory (tenant) ID**

Step 7 Click **Client Credentials**.

Step 8 Unless you already know the client secret *value* (as opposed to the client *secret ID*), you must create a new client secret as follows:

a) Click **New Client Secret**.
b) Enter the required information in the provided fields.
c) Click **Add**.

d) Click **Copy** (⧉) next to Value as the following figure shows.

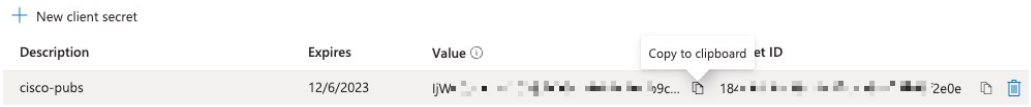

Step 9 From https://portal.azure.com/, click **Event Hubs** > **(name of an event hub)**.

Step 10 In the right pane, click **Copy** (⧉) next to the value of **Host name** and paste the value to the clipboard. This is your event hub host name.

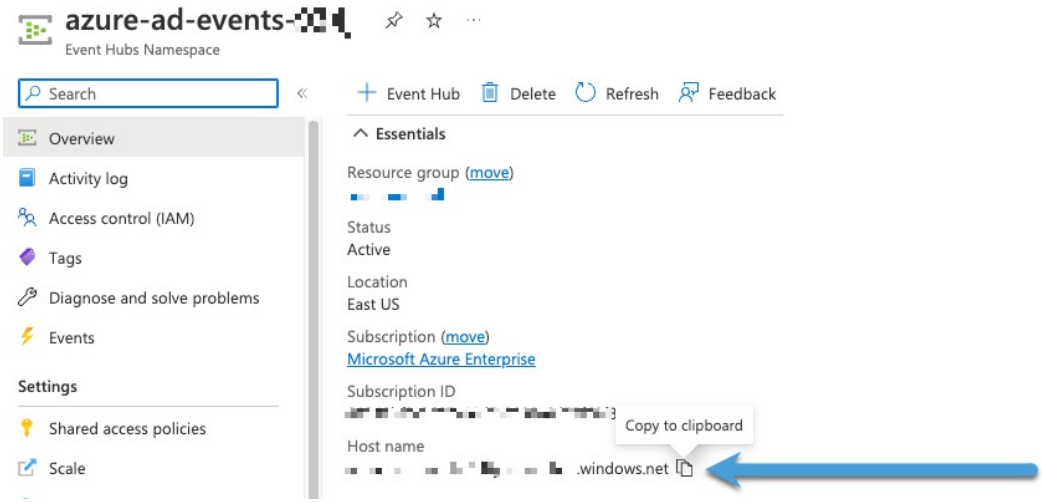

Step 11 Write down or copy to a text file the name of the event hub (same as the **Event Hubs Namespace** at the top of the page).

Step 12 In the left pane, under Settings, click **Shared access policies**.

Step 13 Click the name of a policy.

Step 14 Click **Copy** (⧉) next to **Connection string-primary key**.

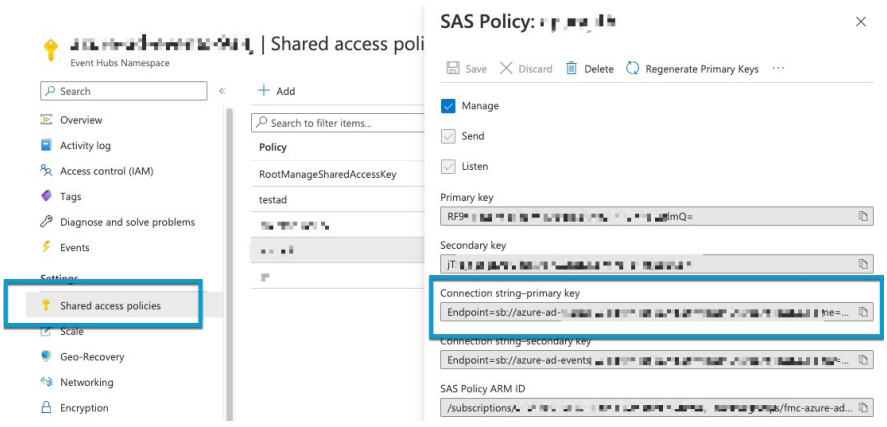

Step 15 Click **Overview** > **Entities** > **Event Hubs** > **(name of an event hub)** > **Entities** > **Consumer Groups**.

Write down the following value or copy it to the clipboard. This is your consumer group name.

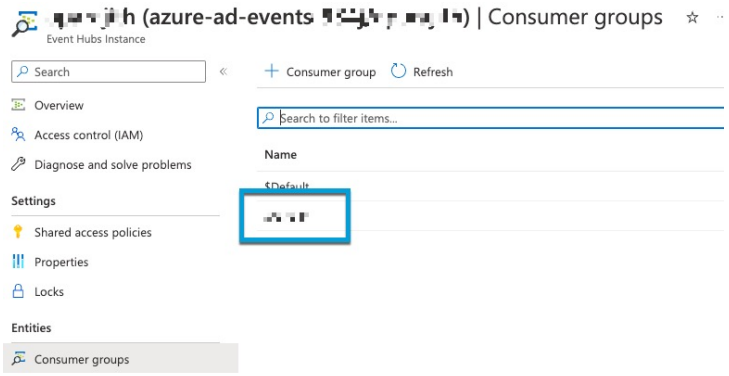

Step 16 In the left pane, click **Overview**.

Step 17 Click **Copy** ( ) next to **Namespace**.

This is your event hubs topic name.

Create an Azure AD Realm

The following procedure enables you to create a *realm* (a connection between the Firewall Management Center and a Microsoft Azure AD realm).

Before you begin

Complete all of the following tasks:

- Configure ISE as discussed in How to Configure ISE for Microsoft Azure AD (SAML), on page 2395
- Create an ISE identity source as discussed in Configure ISE/ISE-PIC, on page 2506
- Enable the Cisco Secure Dynamic Attributes Connector as discussed in Enable the Cisco Secure Dynamic Attributes Connector, on page 1806 .
- Get values required for the Azure AD realm as discussed in Get Required Information For Your Microsoft Azure AD (SAML) Realm, on page 2398.
- Configure Azure AD as discussed in Configure Microsoft Entra ID for Passive Authentication, on page 2395

If you enabled Change Management, you must approve all certificates used in this procedure. Open a new ticket or edit an existing one. For more information, see Creating Change Management Tickets, on page 240 and .Policies and Objects that Support Change Management, on page 235

Note To perform user and identity control with an Azure AD realm, you need only an access control policy with an associated Azure AD realm. You do not need to create an identity policy.

Procedure

Step 1 Log in to Secure Firewall Management Center .

Step 2 Click **Integration** > **Other Integrations** > **Realms** > **Realms**.

Step 3 To create a new realm, click **Add Realm** > **Azure AD**.

Step 4 Enter the following information.

Item	Description
Name	
(Optional.) **Description**	
Client ID	Enter the information you found as discussed in Get Required Information For Your Microsoft Azure AD (SAML) Realm, on page 2398.
Client Secret	
Tenant ID	
Event Hubs Host Name	
Event Hub Name	
Event Hub Connection String	
(Optional.) **Excluded User Groups**	Enter one or more groups from which to *not* download users for identity control. Users in these groups will not be available for use in access control policies. Enter one group name per line followed by a line break. Group names are case-sensitive.

Step 5 To perform other tasks (such as enable, disable, or delete a realm), see Manage a Realm, on page 2455.

Step 6 Enter the values you found as discussed in Get Required Information For Your Microsoft Azure AD (SAML) Realm, on page 2398.

Step 7 Click **Test**.

Step 8 Fix any errors that are displayed in the test.

Step 9 Click **Save**.

What to do next

Create an access control policy and rule as discussed in Creating a Basic Access Control Policy, on page 1748.

Note Before deploying policies related to a Microsoft Azure AD realm, see User Limits for Microsoft Azure Active Directory Realms, on page 2386.

Azure AD (SAML) User Session Timeout

The Azure AD user session timeout for ISE/ISE-PIC is available on the **User Session Timeout** page when you edit an Azure AD realm.

The default value 1440 minutes (24 hours). After the timeout is exceeded, the user's session ends; if the user continues to access the network without logging in again, the user is seen by the Firewall Management Center as Unknown (except for **Failed Captive Portal Users**).

Create an LDAP Realm or an Active Directory Realm and Realm Directory

If you're setting up ISE/ISE-PIC without a realm, be aware there is a user session timeout that affects how users are seen by the Secure Firewall Management Center. For more information, see Realm Fields, on page 2436.

The following procedure enables you to create a *realm* (a connection between the Firewall Management Center and an Active Directory realm) and a *directory* (a connection between the Firewall Management Center and an LDAP server or an Active Directory domain controller).

(Recommended.) To connect securely from the Firewall Management Center to your Active Directory server, first perform the following tasks:

- Export the Active Directory Server's Root Certificate, on page 2443
- Find the Active Directory Server's Name, on page 2442

Microsoft has announced that Active Directory servers will start enforcing LDAP binding and LDAP signing in 2020. Microsoft is making these a requirement because when using default settings, an elevation of privilege vulnerability exists in Microsoft Windows that could allow a man-in-the-middle attacker to successfully forward an authentication request to a Windows LDAP server. For more information, see 2020 LDAP channel binding and LDAP signing requirement for Windows on the Microsoft support site.

For more information about realm and directory configuration fields, see Realm Fields, on page 2436 and Realm Directory and Synchronize fields, on page 2440.

A step-by-step example of setting up a realm with cross-domain trust is shown in Configure the Firewall Management Center for Cross-Domain-Trust: The Setup, on page 2447.

An Active Directory Global Catalog server is *not supported* as a realm directory. For more information about the Global Catalog Server, see Global Catalog on learn.microsoft.com.

Note You must specify a unique **AD Primary Domain** for every Microsoft Active Directory (AD) realm. Although the system allows you to specify the same **AD Primary Domain** for different Microsoft AD realms, the system won't function properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group. The system prevents you from specifying more than one realm with the same **AD Primary Domain** because users and groups won't be identified properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group.

If you're setting up ISE/ISE-PIC without a realm, be aware there is a user session timeout that affects how users are seen by the Secure Firewall Management Center. For more information, see Realm Fields, on page 2436.

Before you begin

If you're using Kerberos authentication for captive portal, see the following section before you begin: Prerequisites for Kerberos Authentication, on page 2436.

If you enabled Change Management, you must open or edit, assign, and approve a ticket for each of the following objects before you can create the realm:

- If you're connecting securely to Microsoft AD or LDAP, the server's trusted certificate authority
- The realm itself

For more information, see Opening a Ticket for Configuration Changes, on page 241 and Policies and Objects that Support Change Management, on page 235.

Important To reduce latency between the Secure Firewall Management Center and your Active Directory domain controller, we strongly recommend you configure a realm directory (that is, domain controller) that is as close as possible geographically to the Secure Firewall Management Center.

For example, if your Secure Firewall Management Center is in North America, configure a realm directory that is also in North America. Failure to do so can cause problems such as timeout downloading users and groups.

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Integration** > **Other Integrations** > **Realms** > **Realms**.

Step 3 To create a new realm, choose from **Add Realm** drop-down list.

Step 4 To perform other tasks (such as enable, disable, or delete a realm), see Manage a Realm, on page 2455.

Step 5 Enter realm information as discussed in Realm Fields, on page 2436.

Step 6 In the Directory Server Configuration section, enter directory information as discussed in Realm Directory and Synchronize fields, on page 2440.

Step 7 (Optional.) To configure another domain for this realm, click **Add another directory**.

Step 8 Click **Configure Groups and Users**.
Enter the following information:

Information	Description
AD Primary Domain	Domain for the Active Directory server where users should be authenticated. For additional information, see Realm Fields, on page 2436. You must create the realm with the original domain name of the domain and not any alternative user principal name (UPN) suffixes of the domain. Otherwise, users and groups fail to download and identity policies will not be enforced. For example, if the original domain is `domain.example.com` and the alternative UPN name is`domain2.mydomain.com`, you must configure the realm to use `domain.example.com`. For more information about configuring an alternative UPN suffix, see a resource like Configuring Alternate Login ID on learn.microsoft.com.
Base DN	The directory tree on the server where the Firewall Management Center should begin searching for user data.
Group DN	The directory tree on the server where the Firewall Management Center should begin searching for group data.
Load Groups	Click to load groups from the Active Directory server. If no groups are displayed, enter or edit information in the **AD Primary Domain**, **Base DN**, and **Group DN** fields and click **Load Groups**. For more information about those fields, see Realm Fields, on page 2436.
Available Groups section	Limit the groups to use in policy by moving them to either the **Included Groups and Users** or **Excluded Groups and Users** list. Moving one group to the **Included Groups and Users** list, for example, allows that group only to be used in policy but excludes all other groups. Groups in the **Excluded Groups and Users** and the users they contain are excluded from excluded from user awareness and control. All other groups and users *are* available. For more information, see Realm Directory and Synchronize fields, on page 2440.

Step 9 Click the **Realm Configuration** tab.

Step 10 Enter **Group Attribute**, and (if you use Kerberos authentication for captive portal) enter **AD Join Username** and **AD Join Password**. For more information, see Realm Directory and Synchronize fields, on page 2440.

Step 11 If you use Kerberos authentication, click **Test**. If the test fails, wait a short time and try again.

Step 12 Enter user session timeout values, in minutes, for **ISE/ISE-PIC Users**, **Terminal Server Agent Users**, **Captive Portal Users**, **Failed Captive Portal Users**, and **Guest Captive Portal Users**.

Step 13 When you're finished configuring the realm, click **Save**.

What to do next

- Configure the Firewall Management Center for Cross-Domain-Trust: The Setup, on page 2447

- Synchronize Users and Groups, on page 2445
- Edit, delete, enable, or disable a realm; see Manage a Realm, on page 2455.
- Compare Realms, on page 2456.
- Optionally, monitor the task status; see *Viewing Task Messages* in the Cisco Secure Firewall Management Center Administration Guide.

About Realms and Realm Sequences

Realms are connections between the Secure Firewall Management Center and the user accounts on the servers you monitor. They specify the connection settings and authentication filter settings for the server. Realms can:

- Specify the users and user groups whose activity you want to monitor.
- Query the user repository for user metadata on authoritative users, as well as some non-authoritative users: POP3 and IMAP users detected by traffic-based detection and users detected by traffic-based detection, a TS Agent, or ISE/ISE-PIC.

(Microsoft AD realm only.) A *realm sequence* is an ordered list of two or more Active Directory realms to use in identity policy. When you associate a realm sequence with an identity rule, the system searches the Active Directory domains in order from first to last as specified in the realm sequence.

You can add multiple domain controllers as directories in a realm, but they must share the same basic realm information. The directories in a realm must be exclusively LDAP or exclusively Active Directory (AD) servers. After you enable a realm, your saved changes take effect next time the Secure Firewall Management Center queries the server.

To perform user awareness, you must configure a realm for any of the supported server types. The system uses these connections to query the servers for data associated with POP3 and IMAP users, and to collect data about LDAP users discovered through traffic-based detection.

The system uses the email addresses in POP3 and IMAP logins to correlate with LDAP users on an Active Directory, Microsoft Azure Active Directory, or OpenLDAP. For example, if a managed device detects a POP3 login for a user with the same email address as an LDAP user, the system associates the LDAP user's metadata with that user.

To perform user control, you can configure any of the following:

- A realm or realm sequence for an Active Directory, Microsoft Azure Active Directory, server, or for ISE/ISE-PIC

Note Configuring a Microsoft AD realm or realm sequence is optional if you plan to configure SGT ISE attribute conditions but not user, group, realm, Endpoint Location, or Endpoint Profile conditions; or if you use your identity policy only to filter network traffic.

Realm sequence is not available for Microsoft Azure AD realms.

- A realm or realm sequence for a Microsoft AD server for the TS Agent.
- For captive portal, an LDAP realm.

A realm sequence is not supported for LDAP.

You can nest Microsoft AD groups and the Secure Firewall Management Center downloads those groups and the users they contain. You can optionally restrict which groups and users get downloaded as discussed in Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.

About User Synchronization

You can configure a realm or realm sequence to establish a connection between the Firewall Management Center and an LDAP or Microsoft AD server to retrieve user and user group metadata for certain detected users:

- LDAP and Microsoft AD users authenticated by captive portal or reported by ISE/ISE-PIC. This metadata can be used for user awareness and user control.
- POP3 and IMAP user logins detected by traffic-based detection, if those users have the same email address as an LDAP or AD user. This metadata can be used for user awareness.

The Firewall Management Center obtains the following information and metadata about each user:

- LDAP user name
- First and last names
- Email address
- Department
- Telephone number

Important To reduce latency between the Secure Firewall Management Center and your Active Directory domain controller, we strongly recommend you configure a realm directory (that is, domain controller) that is as close as possible geographically to the Secure Firewall Management Center.

For example, if your Secure Firewall Management Center is in North America, configure a realm directory that is also in North America. Failure to do so can cause problems such as timeout downloading users and groups.

About User Activity Data

User activity data is stored in the user activity database and user identity data is stored in the users database. The maximum number of users you can store and use in access control depends on your Firewall Management Center model. When choosing which users and groups to include, make sure the total number of users is less than your model limit. If your access control parameters are too broad, the Firewall Management Center obtains information on as many users as it can and reports the number of users it failed to retrieve in the Tasks tab page of the Message Center.

To optionally limit the subnets on which a managed device watches for user awareness data, you can use the **configure identity-subnet-filter** command as discussed in the Cisco Secure Firewall Threat Defense Command Reference.

Note If you remove a user that has been detected by the system from your user repository, the Firewall Management Center does *not* remove that user from its users database; you must manually delete it. However, your LDAP changes *are* reflected in access control rules when the Firewall Management Center next updates its list of authoritative users.

Realms and Trusted Domains

When you configure a Microsoft Active Directory (AD) *realm* in the Firewall Management Center, it is associated with a Microsoft Active Directory or LDAP *domain*.

A grouping of Microsoft Active Directory (AD) domains that trust each other is commonly referred to as a *forest*. This trust relationship can enable domains to access each other's resources in different ways. For example, a user account defined in domain A can be marked as a member of a group defined in domain B.

Note Trusted domains apply to Microsoft Active Directory domains only. They do *not* apply to either Microsoft Azure Active Directory or LDAP domains.

The system and trusted domains

The system supports AD forests that are configured in a trust relationship. There are several types of trust relationships; this guide discusses two-way, transitive forest trust relationships. The following simple example shows two forests: **forest.example.com** and **eastforest.example.com**. Users and groups in each forest can be authenticated by AD in the other forest, provided you configure the forests that way.

If you set up the system with one realm for each domain and one directory for each domain controller, the system can discover up to 100,000 foreign security principals (users and groups). If these foreign security principals match a user downloaded in another realm, then they can be used in access control policy.

You need not configure a realm for any domain that has no users you wish to use in access control policies.

To continue the example, suppose you have three AD forests (one of which could be a subdomain or an independent forest), all set up as two-way transitive forest relationships, all users and groups are available in all three forests as well as in the system. (As in the preceding example, all three AD domains must be set up as realms and all domain controllers must be configured as directories in those realms.)

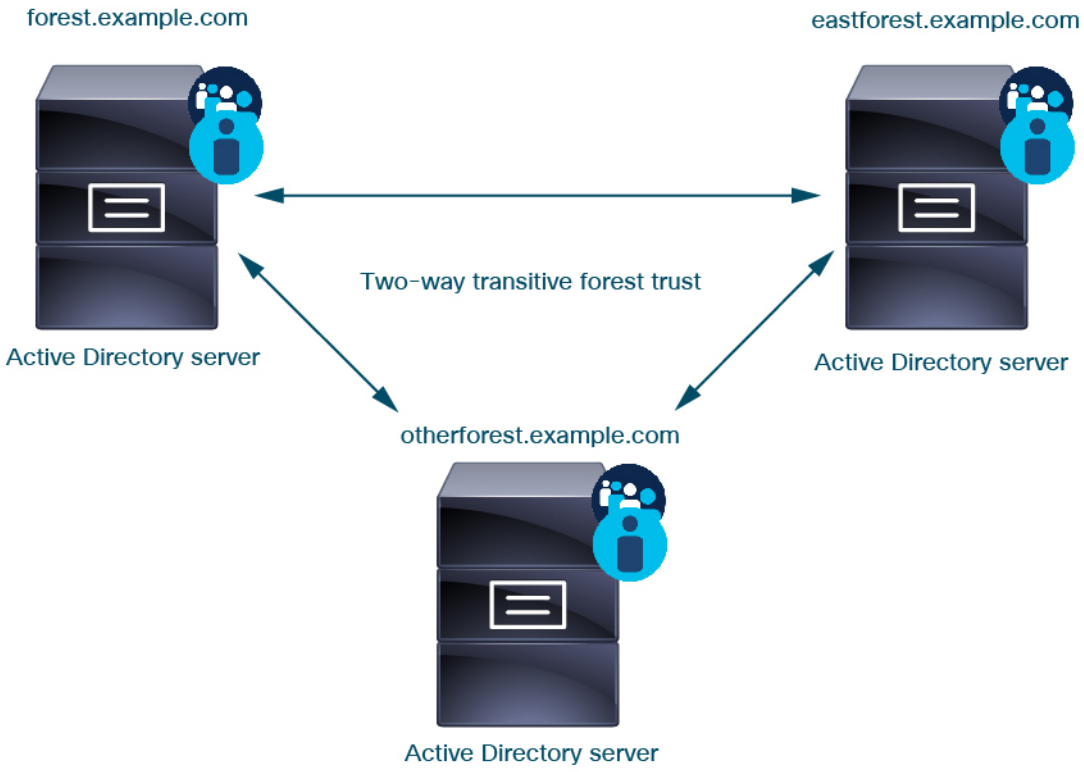

Finally, you can set up the Firewall Management Center to be able to enforce identity policies on users and groups in a two-forest system with two-way transitive forest trust. Suppose each forest has at least one domain controller, each of which authenticates different users and groups. For the Firewall Management Center to be able to enforce identity policies on those users and groups, you must set up each domain containing relevant users as Firewall Management Center realm and each domain controller as Firewall Management Center directory in the respective realm.

Failure to properly configure the Firewall Management Center prevents some of the users and groups from being able to be used in policies. You will see warnings when you try to synchronize users and groups in that case.

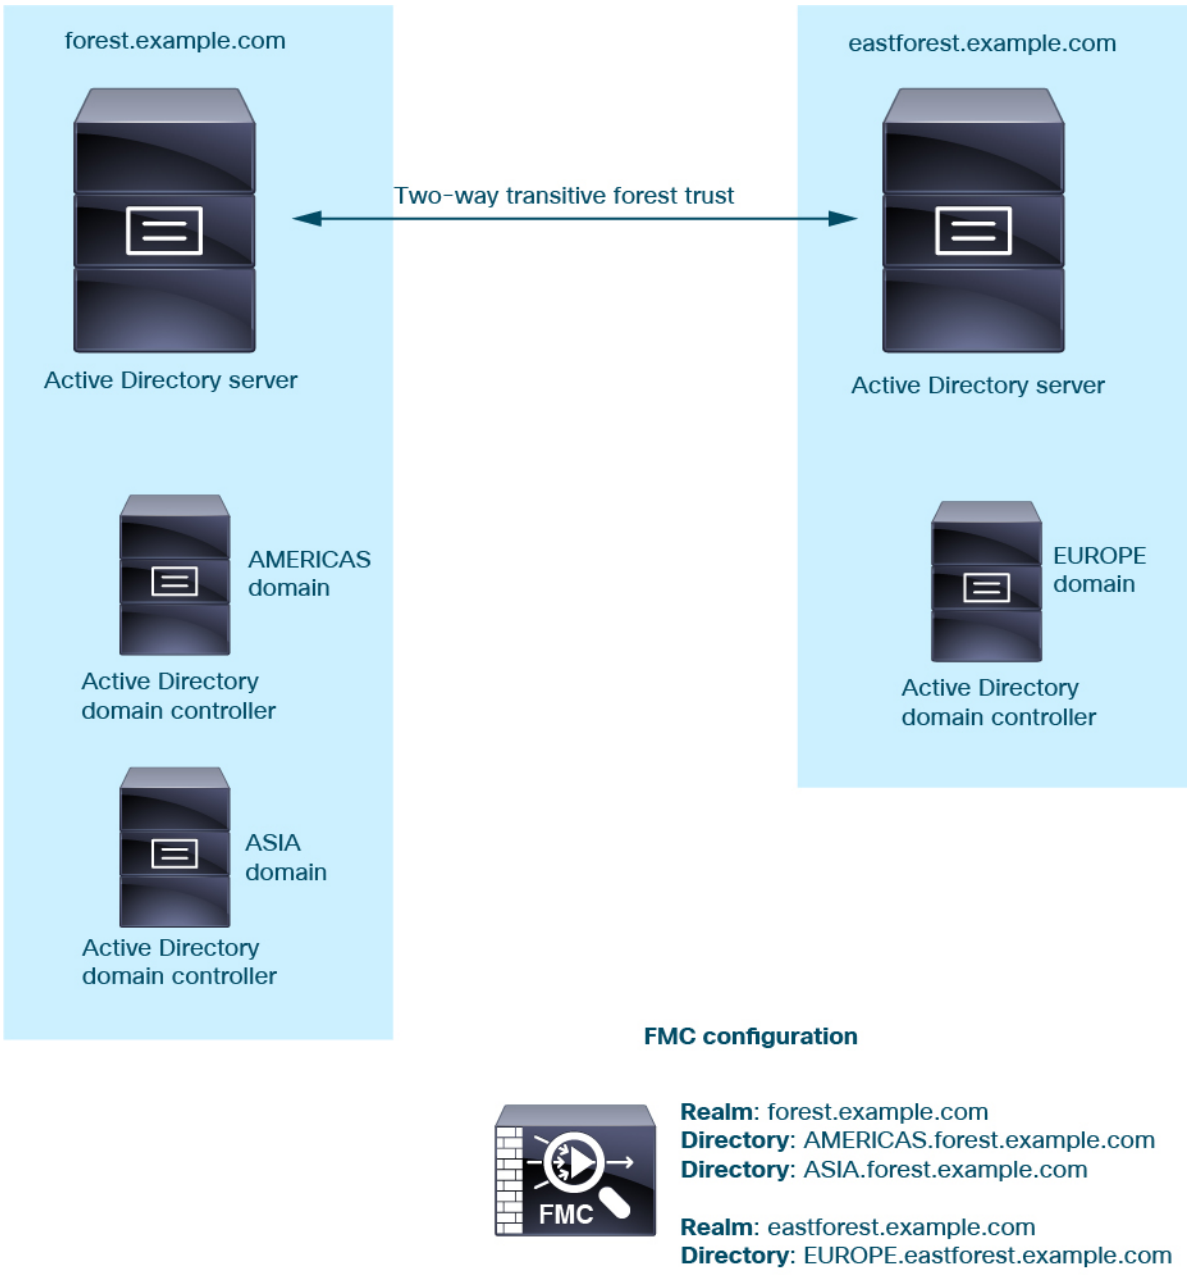

Using the preceding example, set up the Firewall Management Center as follows:

- Realm for any domain in **forest.example.com** that contains users you want to control with access control policies
  - Directory in the realm for **AMERICAS.forest.example.com**
  - Directory in the realm for **ASIA.forest.example.com**
- Realm for any domain in **eastforest.example.com** that contains users you want to control with access control policies

- Directory in the realm for **EUROPE.eastforest.example.com**

Note The Firewall Management Center uses the AD field **msDS-PrincipalName** to resolve references to find user and group names in each domain controller. **msDS-PrincipalName** returns a NetBIOS name.

Supported Servers for Realms

You can configure realms to connect to the following types of servers, providing they have TCP/IP access from the Firewall Management Center:

Server Type	Supported for ISE/ISE-PIC data retrieval?	Supported for TS Agent data retrieval?	Supported for captive portal data retrieval?
Microsoft Active Directory on Windows Server 2012, 2016, and 2019	Yes	Yes	Yes
Microsoft Azure AD	Yes	No	No
OpenLDAP on Linux	No	No	Yes

An Active Directory Global Catalog server is *not supported* as a realm directory. For more information about the Global Catalog Server, see Global Catalog on learn.microsoft.com.

Note If the TS Agent is installed on a Microsoft Active Directory Windows Server shared with another passive authentication identity source (ISE/ISE-PIC), the Firewall Management Center prioritizes the TS Agent data. If the TS Agent and a passive identity source report activity by the same IP address, only the TS Agent data is logged to the Firewall Management Center.

Note the following about your server group configurations:

- To perform user control on user groups or on users in groups, you must configure user groups on the LDAP or Active Directory server.
- Group names cannot start with `s-` because it is used internally by LDAP.

  Neither group names nor organizational unit names can contain special characters like asterisk (*), equals (=), or backslash (\); otherwise, users in those groups or organizational units are not downloaded and are not available for identity policies.
- To configure an Active Directory realm that includes or excludes users who are members of a sub-group on your server, note that Microsoft recommends that Active Directory has no more than 5000 users per group in Windows Server 2012. For more information, see Active Directory Maximum Limits—Scalability on MSDN.

If necessary, you can modify your Active Directory server configuration to increase this default limit and accommodate more users.

- To uniquely identify the users reported by a server in your Remote Desktop Services environment, you must configure the Cisco Terminal Services (TS) Agent. When installed and configured, the TS Agent assigns unique ports to individual users so the system can uniquely identify those users. (Microsoft changed the name *Terminal Services* to *Remote Desktop Services*.)

  For more information about the TS Agent, see the *Cisco Terminal Services (TS) Agent Guide*.

Supported Server Object Class and Attribute Names

The servers in your realms *must* use the attribute names listed in the following table for the Secure Firewall Management Center to retrieve user metadata from the servers. If the attribute names are incorrect on your server, the Secure Firewall Management Center cannot populate its database with the information in that attribute.

Table 135: Map of attribute names to Secure Firewall Management Center fields

Metadata	Firewall Management Center Attribute	LDAP ObjectClass	Active Directory Attribute	OpenLDAP Attribute
LDAP user name	Username	• user • inetOrgPerson	samaccountname	cn uid
first name	First Name		givenname	givenname
last name	Last Name		sn	sn
email address	Email		mail userprincipalname (if mail has no value)	mail
department	Department		department distinguishedname (if department has no value)	ou
telephone number	Phone		telephonenumber	telephonenumber

Note The LDAP ObjectClass for groups is `group`, `groupOfNames`, (`group-of-names` for Active Directory) or `groupOfUniqueNames`.

For more information about ObjectClasses and attributes, see the following references:

- Microsoft Active Directory:
  - ObjectClasses: All Classes on MSDN

- Attributes: All Attributes on MSDN

- OpenLDAP: RFC 4512

Prerequisites for Kerberos Authentication

If you're using Kerberos to authentication captive portal users, keep the following in mind.

Hostname character limit

If you're using Kerberos authentication, the managed device's host name must be less than 15 characters (it's a NetBIOS limitation set by Windows); otherwise, captive portal authentication fails. You set the managed device host name when you set up the device. For more information, see an article like this one on the Microsoft documentation site: Naming conventions in Active Directory for computers, domains, sites, and OUs.

DNS response character limit

DNS must return a response of 64KB or less to the hostname; otherwise, the AD connection test fails. This limit applies in both directions and is discussed in RFC 6891 section-6.2.5.

Realm Fields

The following fields are used to configure a realm.

Realm Configuration Fields

These settings apply to all Active Directory servers or domain controllers (also referred to as *directories*) in a realm.

Name

A unique name for the realm.

- To use the realm in identity policies, the system supports alphanumeric and special characters.
- To use the realm in RA VPN configurations, the system supports alphanumeric, hyphen (-), underscore (_), and plus (+) characters.

Description

(Optional.) Enter a description of the realm.

Type

The type of realm, **AD** for Microsoft Active Directory, **LDAP** for other supported LDAP repositories, or **Local**. For a list of supported LDAP repositories, see Supported Servers for Realms, on page 2434. You can authenticate captive portal users with an LDAP repository; all others require Active Directory.

Note Only captive portal supports an LDAP realm.

The realm type **LOCAL** is used for configuring local user settings. The LOCAL realm is used in remote access user authentication.

Add the following Local User Information for the LOCAL realm:

- **Username**—Name of the local user.
- **Password**—Local user password.
- **Confirm Password**—Confirm the local user password.

Note **Click Add another local user** to add more users to the LOCAL realm.

You can add more users after creating the realm and update password for the local users. You can also create multiple LOCAL realms but cannot disable them.

AD Primary Domain

For Microsoft Active Directory realms only. Domain for the Active Directory server where users should be authenticated.

Note You must specify a unique **AD Primary Domain** for every Microsoft Active Directory (AD) realm. Although the system allows you to specify the same **AD Primary Domain** for different Microsoft AD realms, the system won't function properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group. The system prevents you from specifying more than one realm with the same **AD Primary Domain** because users and groups won't be identified properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group.

AD Join Username and AD Join Password
(Available on the **Realm Configuration** tab page when you edit a realm.)

For Microsoft Active Directory realms intended for Kerberos captive portal active authentication, the distinguished username and password of any Active Directory user with appropriate rights to create a Domain Computer account in the Active Directory domain.

Keep the following in mind:

- DNS must be able to resolve the domain name to an Active Directory domain controller's IP address.
- The user you specify must be able to join computers to the Active Directory domain.
- The user name must be fully qualified (for example, **`administrator@mydomain.com`**, *not* **`administrator`**).

If you choose **`Kerberos`** (or **`HTTP Negotiate`**, if you want Kerberos as an option) as the **Authentication Protocol** in an identity rule, the **Realm** you select must be configured with an **AD Join Username** and **AD Join Password** to perform Kerberos captive portal active authentication.

Note The SHA-1 hash algorithm is not secure for storing passwords on your Active Directory server and should not be used. For more information, consult a reference such as Migrating your Certification Authority Hashing Algorithm from SHA1 to SHA2 on Microsoft TechNet or Password Storage Cheat Sheet on the Open Web Application Security Project website.

We recommend SHA-256 for communicating with Active Directory.

Directory Username and Directory Password

The distinguished username and password for a user with appropriate access to the user information you want to retrieve.

Note the following:

- For some versions of Microsoft Active Directory, specific permissions might be required to read users and groups. Consult the documentation provided with Microsoft Active Directory for details.
- For OpenLDAP, the user's access privileges are determined by the `<level>` parameter discussed in section 8 of the OpenLDAP specification. The user's `<level>` should be `auth` or better.
- The user name must be fully qualified (for example, **`administrator@mydomain.com`**, *not* **`administrator`**).

Note The SHA-1 hash algorithm is not secure for storing passwords on your Active Directory server and should not be used. For more information, consult a reference such as Migrating your Certification Authority Hashing Algorithm from SHA1 to SHA2 on Microsoft TechNet or Password Storage Cheat Sheet on the Open Web Application Security Project website.

We recommend SHA-256 for communicating with Active Directory.

Base DN

(Optional.) The directory tree on the server where the Secure Firewall Management Center should begin searching for user data. If you don't specify a **Base DN**, the system retrieves the top-level DN provided you can connect to the server.

Typically, the base distinguished name (DN) has a basic structure indicating the company domain name and operational unit. For example, the Security organization of the Example company might have a base DN of **`ou=security,dc=example,dc=com`**.

Group DN

(Optional.) The directory tree on the server where the Secure Firewall Management Center should search for users with the group attribute. A list of supported group attributes is shown in Supported Server Object Class and Attribute Names, on page 2435. If you don't specify a **Group DN**, the system retrieves the top-level DN provided you can connect to the server.

Note Following is the list of characters the system *supports* in users, groups, DNs in your directory server. Using any characters other than the following could result in the system failing to download users and groups.

Entity	Supported characters
User name	**a-z A-Z 0-9 ! # $ % ^ & ( ) _ - { } ' . ~ `**
Group name	**a-z A-Z 0-9 ! # $ % ^ & ( ) _ - { } ' . ~ `**
Base DN and Group DN	**a-z A-Z 0-9 ! @ $ % ^ & * ( ) _ - . ~ `**

A space is not supported anywhere in a user name, including at the end.

The following fields are available when you edit an existing realm.

User Session Timeout

(Available on the **Realm Configuration** tab page when you edit a realm.)

Enter the number of minutes before user sessions time out. The default is 1440 (24 hours) after the user's login event. After the timeout is exceeded, the user's session ends; if the user continues to access the network without logging in again, the user is seen by the Secure Firewall Management Center as Unknown (except for **Failed Captive Portal Users**).

In addition, if you set up ISE/ISE-PIC without a realm and the timeout is exceeded, a workaround is required. For more information, contact Cisco TAC.

You can set timeout values for the following:

- **User Agent and ISE/ISE-PIC Users**: Timeout for users tracked by the user agent or by ISE/ISE-PIC, which are types of passive authentication.

  The timeout value you specify does *not* apply to pxGrid SXP session topic subscriptions (for example, destination SGT mappings). Instead, session topic mappings are preserved as long as there is no delete or update message for a given mapping from ISE.

  For more information about ISE/ISE-PIC, see The ISE/ISE-PIC Identity Source, on page 2499.

- **Terminal Services Agent Users**: Timeout for users tracked by the TS Agent, which is a type of passive authentication. For more information, see The Terminal Services (TS) Agent Identity Source, on page 2549.

- **Captive Portal Users**: Timeout for users who successfully log in using the captive portal, which is a type of active authentication. For more information, see The Captive Portal Identity Source, on page 2525.

- **Failed Captive Portal Users**: Timeout for users who do not successfully log in using the captive portal. You can configure the **Maximum login attempts** before the user is seen by the Secure Firewall Management Center as Failed Auth User. A Failed Auth User can optionally be granted access to the network using access control policy and, if so, this timeout value applies to those users.

  For more information about failed captive portal logins, see Captive Portal Fields, on page 2539.

- **Guest Captive Portal Users**: Timeout for users who log in to the captive portal as a guest user. For more information, see The Captive Portal Identity Source, on page 2525.

Realm Directory and Synchronize fields

Realm Directory Fields

These settings apply to individual servers (such as Active Directory domain controllers) in a realm.

Hostname / IP Address

Fully qualified host name of the Active Directory domain controller machine. To find the fully qualified name, see Find the Active Directory Server's Name, on page 2442.

If you're using Kerberos for authenticating captive portal, also make sure you understand the following:

If you're using Kerberos authentication, the managed device's host name must be less than 15 characters (it's a NetBIOS limitation set by Windows); otherwise, captive portal authentication fails. You set the managed device host name when you set up the device. For more information, see an article like this one on the Microsoft documentation site: Naming conventions in Active Directory for computers, domains, sites, and OUs.

DNS must return a response of 64KB or less to the hostname; otherwise, the AD connection test fails. This limit applies in both directions and is discussed in RFC 6891 section-6.2.5.

Port

The server's port.

Encryption

(Strongly recommended.) The encryption method to use:

- **`STARTTLS`**—encrypted LDAP connection
- **`LDAPS`**—encrypted LDAP connection
- **`None`**—unencrypted LDAP connection (unsecured traffic)

To communicate securely with an Active Directory server, see Connect Securely to Active Directory or LDAP, on page 2442.

CA Certificate

The TLS/SSL certificate to use for authentication to the server. You must configure **`STARTTLS`** or **`LDAPS`** as the **Encryption** type to use a TLS/SSL certificate.

If you are using a certificate to authenticate, the name of the server in the certificate must match the server **Hostname / IP Address**. For example, if you use 10.10.10.250 as the IP address but **`computer1.example.com`** in the certificate, the connection fails.

Interface used to connect to Directory server

Required only for RA VPN authentication so the Secure Firewall Threat Defense can connect securely to your Active Directory server. This interface is not used for downloading users and groups, however.

You can choose only a routed interface group. For more information, see Interface, on page 1398.

Click one of the following:

- **Resolve via route lookup**: Use routing to connect to the Active Directory server.
- **Choose an interface**: Choose a specific managed device interface group to connect to the Active Directory server.

User Synchronize Fields

AD Primary Domain

For Microsoft Active Directory realms only. Domain for the Active Directory server where users should be authenticated.

Note You must specify a unique **AD Primary Domain** for every Microsoft Active Directory (AD) realm. Although the system allows you to specify the same **AD Primary Domain** for different Microsoft AD realms, the system won't function properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group. The system prevents you from specifying more than one realm with the same **AD Primary Domain** because users and groups won't be identified properly. This happens because system assigns a unique ID to every user and group in each *realm*; therefore, the system cannot definitively identify any particular user or group.

Enter query to look for users and groups

Base DN:

(Optional.) The directory tree on the server where the Firewall Management Center should begin searching for user data.

Typically, the base distinguished name (DN) has a basic structure indicating the company domain name and operational unit. For example, the Security organization of the Example company might have a base DN of `ou=security,dc=example,dc=com`.

Group DN:

(Optional.) The directory tree on the server where the Firewall Management Center should search for users with the group attribute. A list of supported group attributes is shown in Supported Server Object Class and Attribute Names, on page 2435.

Note Neither the group name nor the organizational unit name can contain special characters like asterisk (*), equals (=), backslash (\) because users in those groups are not downloaded and cannot be used in identity policies.

Load Groups

Enables you to download users and groups for user awareness and user control.

Available Groups, Add to Include, Add to Exclude

Limits the groups that can be used in policy.

- Groups that are displayed in the **Available Groups** field are available for policy unless you move groups to the **Included Groups and Users** or **Excluded Groups and Users** field.
- If you move groups to the **Included Groups and Users** field, only those groups and users they contain are downloaded and user data is available for user awareness and user control.
- If you move groups to the **Excluded Groups and Users** field, all groups and users they contain *except* these are downloaded and available for user awareness and user control.

- To include users from groups that are not included, enter the user name in the field below **User Inclusion** and click **Add**.
- To exclude users from groups that are not excluded, enter the user name in the field below **User Exclusion** and click **Add**.

Note The users that are downloaded to the Firewall Management Center is calculated using the formula **R = I - (E+e) + i**, where

- R is list of downloaded users
- I is included groups
- E is excluded groups
- e is excluded users
- i is included users

Synchronize Now
Click to synchronize groups and users with AD.

Begin automatic synchronization at
Enter the time and time interval at which to download users and groups from AD.

Connect Securely to Active Directory or LDAP

To create a secure connection between an Active Directory or LDAP server and the Secure Firewall Management Center (which we strongly recommend), you must perform all of the following tasks:

- Export the server's *root* certificate.
- Import the root certificate into the Secure Firewall Management Center as a trusted CA certificate (**Objects** > **Object Management** > **PKI** > **Trusted CAs**).
- Find the server's fully qualified name.
- Create the realm directory.

See one of the following tasks for more information.

Related Topics

Export the Active Directory Server's Root Certificate, on page 2443
Find the Active Directory Server's Name, on page 2442
Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426

Find the Active Directory Server's Name

To configure a realm directory in the Firewall Management Center, you must know the fully qualified server name, which you can find as discussed in the procedure that follows.

These tasks apply to Microsoft Active Directory only. If you use LDAP, consult an appropriate reference for the procedure.

Before you begin

You must log in to the Active Directory server as a user with sufficient privileges to view the computer's name.

Procedure

Step 1 Log in to the Active Directory server.

Step 2 Click **Start**.

Step 3 Right-click **This PC**.

Step 4 Click **Properties**.

Step 5 Click **Advanced System Settings**.

Step 6 Click the **Computer Name** tab.

Step 7 Note the value of **Full computer name**.
You must enter this exact name when you configure the realm directory in the Firewall Management Center.

What to do next

Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.

Related Topics

Export the Active Directory Server's Root Certificate, on page 2443

Export the Active Directory Server's Root Certificate

The task that follows discusses how to export the Active Directory server's root certificate, which is required to connect securely to the Firewall Management Center to obtain user identity information.

These tasks apply to Microsoft Active Directory only. If you use LDAP, consult an appropriate reference for the procedure.

Before you begin

You must know the name of your Active Directory server's root certificate. The root certificate might have the same name as the domain or the certificate might have a different name. The procedure that follows shows one way you can find the name; there could be other ways, however.

Procedure

Step 1 Following is one way to find the name of the Active Directory Server's root certificate; consult Microsoft documentation for more information:

a) Log in to the Active Directory server as a user with privileges to run the Microsoft Management Console.
b) Click **Start** and enter **mmc**.
c) Click **File** > **Add/Remove Snap-in**
d) From the Available Snap-ins list in the left pane, click **Certificates (local)**.
e) Click **Add**.

f) At the Certificates snap-in dialog box, click **Computer Account** and click **Next**.
g) At the Select Computer dialog box, click **Local Computer** and click **Finish**.
h) *Windows Server 2012 only*. Repeat the preceding steps to add the Certification Authority snap-in.
i) Click **Console Root** > **Trusted Certification Authorities** > **Certificates**.
The server's trusted certificates are displayed in the right pane. The following figure is only an example for Windows Server 2012; yours will probably look different.

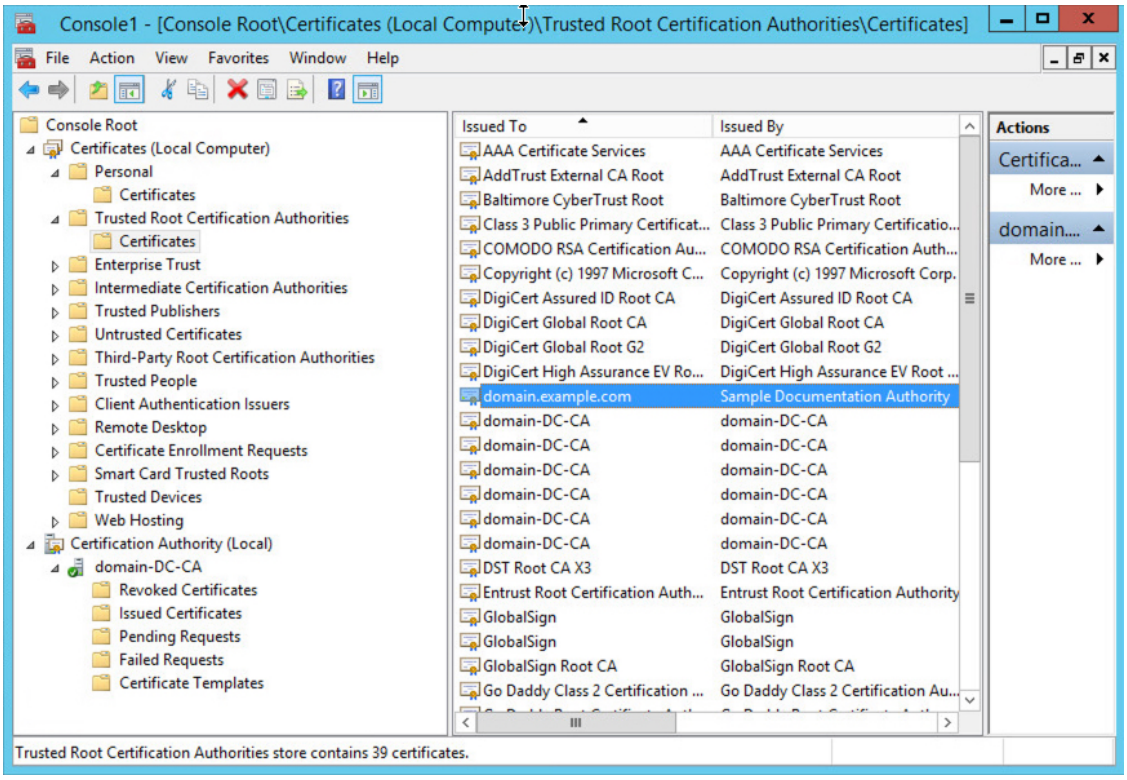

Step 2 Export the certificate using the **certutil** command.

This is only one way to export the certificate. It's a convenient way to export the certificate, especially if you can run a web browser and connect to the Firewall Management Center from the Active Directory server.

a) Click **Start** and enter **cmd**.
b) Enter the command **certutil -ca.cert** *certificate-name*.
The server's certificate is displayed on the screen.
c) Copy the entire certificate to the clipboard, starting with **-----BEGIN CERTIFICATE-----** and ending with **-----END CERTIFICATE-----** (including those strings).

What to do next

Import the Active Directory server's certificate into the Firewall Management Center as a Trusted CA Certificate as discussed in Adding a Trusted CA Object, on page 1409.

Related Topics

Find the Active Directory Server's Name, on page 2442

Synchronize Users and Groups

Synchronizing users and groups means the Firewall Management Center queries the realms and directories you configured for groups and users in those groups. All users the Firewall Management Center finds can be used in identity policies.

If issues are found, you most likely need to add a realm that contains users and groups the Firewall Management Center cannot load. For details, see Realms and Trusted Domains, on page 2431.

Before you begin

Create a Secure Firewall Management Center *realm* for each Active Directory domain and a Firewall Management Center *directory* for each Active Director domain controller in each forest. See Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.

Note Synchronizing users and groups is not necessary for Microsoft Azure AD realms.

You must create a realm only for domains that have users you want to use in user control.

You can nest Microsoft AD groups and the Secure Firewall Management Center downloads those groups and the users they contain. You can optionally restrict which groups and users get downloaded as discussed in Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.

You must create the realm with the original domain name of the domain and not any alternative user principal name (UPN) suffixes of the domain. Otherwise, users and groups fail to download and identity policies will not be enforced. For example, if the original domain is `domain.example.com` and the alternative UPN name is`domain2.mydomain.com`, you must configure the realm to use `domain.example.com`. For more information about configuring an alternative UPN suffix, see a resource like Configuring Alternate Login ID on learn.microsoft.com.

Procedure

Step 1 Log in to the Secure Firewall Management Center if you have not already done so.

Step 2 Click **Integration** > **Other Integrations** > **Realms** > **Realms**.

Step 3 Next to each realm, click **Download** ( ).

Step 4 To see the results, click the **Sync Results** tab.
The Realms column indicates whether or not there were issues synchronizing users and groups in Active Directory forests. Look for the following indicators next to each realm.

Indicator in Realms column	Meaning
(nothing)	All users and groups synchronized without error. No action is necessary.

Indicator in Realms column	Meaning
Yellow Triangle (⚠)	There were issues synchronizing users and groups. Make sure you added a realm for each Active Directory domain and a directory for each Active Directory domain controller. For more details, see Troubleshoot Cross-Domain Trust, on page 2461.

Create a Realm Sequence

The following procedure enables you to create a realm sequence, which is an ordered list of realms the system searches when it applies identity policy. You add a realm sequence to an identity rule exactly the same way as you add a realm; the difference is that the system searches all the realms in the order specified in the realm sequence when applying an identity policy.

Before you begin

You must create and enable at least two realms, each corresponding to a connection with an Active Directory server. You cannot create realm sequences for LDAP realms.

Create a realm as discussed in Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.

Procedure

Step 1 Log in to the Secure Firewall Management Center if you have not already done so.

Step 2 Click **Integration** > **Other Integrations** > **Realms** > **Realm Sequences**.

Step 3 Click the **Realm Sequences** tab.

Step 4 Click **Add Sequence**.

Step 5 In the **Name** field, enter a name to identify the realm sequence.

Step 6 (Optional.) In the **Description** field, enter a description for the realm sequence.

Step 7 Under Realms, click **Add** (+).

Step 8 Click the name of each realm to add to the sequence.

To narrow your search, enter all or part of a realm name in **Filter** field.

Step 9 Click **OK**.

Step 10 In the Add Realm Sequence dialog box, drag and drop the realms in the order in which you want the system to search for them.
The following figure shows an example of a realm sequence consisting of two realms. The **domain-europe.example.com** realm will be searched for users before the **domain.example.com** realm.

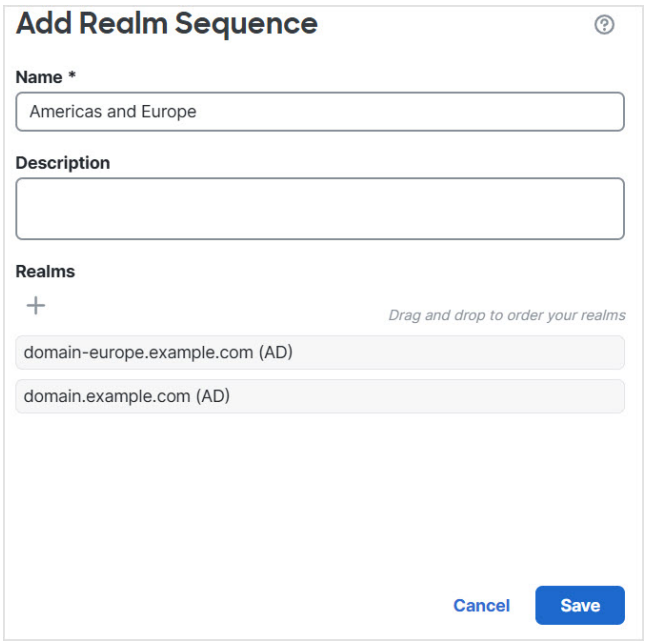

Step 11 Click **Save**.

What to do next

See Create an Identity Policy, on page 2555.

Configure the Firewall Management Center for Cross-Domain-Trust: The Setup

This is an introduction to several topics that walk you through configuring the Firewall Management Center with two realms with cross-domain trust.

This step-by-step example involves two forests: **forest.example.com** and **eastforest.example.com**. The forests are configured so that certain users and groups in each forest can be authenticated by Microsoft AD in the other forest.

Note This topic applies to Microsoft AD realms only. It does *not* apply to Microsoft Azure AD realms.

Following is the example setup used in this example.

Using the preceding example, you would set up the Firewall Management Center as follows:

- Realm and directory for any domain in **forest.example.com** that contains users you want to control with access control policy
- Realm and directory for any domain in **eastforest.example.com** that contains users you want to control with access control policy

Each realm in the example has one domain controller, which is configured in the Firewall Management Center as a directory. The directories in this example are configured as follows:

- **forest.example.com**
  - Base distinguished name (DN) for users: **ou=UsersWest,dc=forest,dc=example,dc=com**
  - Base DN for groups: **ou=EngineringWest,dc=forest,dc=example,dc=com**
- **eastforest.example.com**
  - Base DN for users: **ou=EastUsers,dc=eastforest,dc=example,dc=com**
  - Base DN for groups: **ou=EastEngineering,dc=eastforest,dc=example,dc=com**

Related Topics

Configure the Secure Firewall Management Center for Cross-Domain-Trust Step 1: Configure Realms and Directories, on page 2448

Configure the Secure Firewall Management Center for Cross-Domain-Trust Step 1: Configure Realms and Directories

This is the first task in a step-by-step procedure that explains how to configure the Firewall Management Center to recognize Active Directory servers configured in a cross-domain trust relationship, which is an increasingly common configuration for enterprise organizations. For an overview of this sample configuration, see Configure the Firewall Management Center for Cross-Domain-Trust: The Setup, on page 2447.

If you set up the system with one realm for each domain and one directory for each domain controller, the system can discover up to 100,000 foreign security principals (users and groups). If these foreign security principals match a user downloaded in another realm, then they can be used in access control policy.

Before you begin

You must configure Microsoft Active Directory servers in a cross-domain trust relationship; see Realms and Trusted Domains, on page 2431 for more information.

If you authenticate users with LDAP or Microsoft Azure AD, you *cannot* use this procedure.

Procedure

Step 1 Log in to the Secure Firewall Management Center if you have not already done so.

Step 2 Click **Integration** > **Other Integrations** > **Realms** > **Realms**.

Step 3 Click **Add Realm** > **Active Directory/LDAP**.

Step 4 Enter the following information to configure **forest.example.com**.

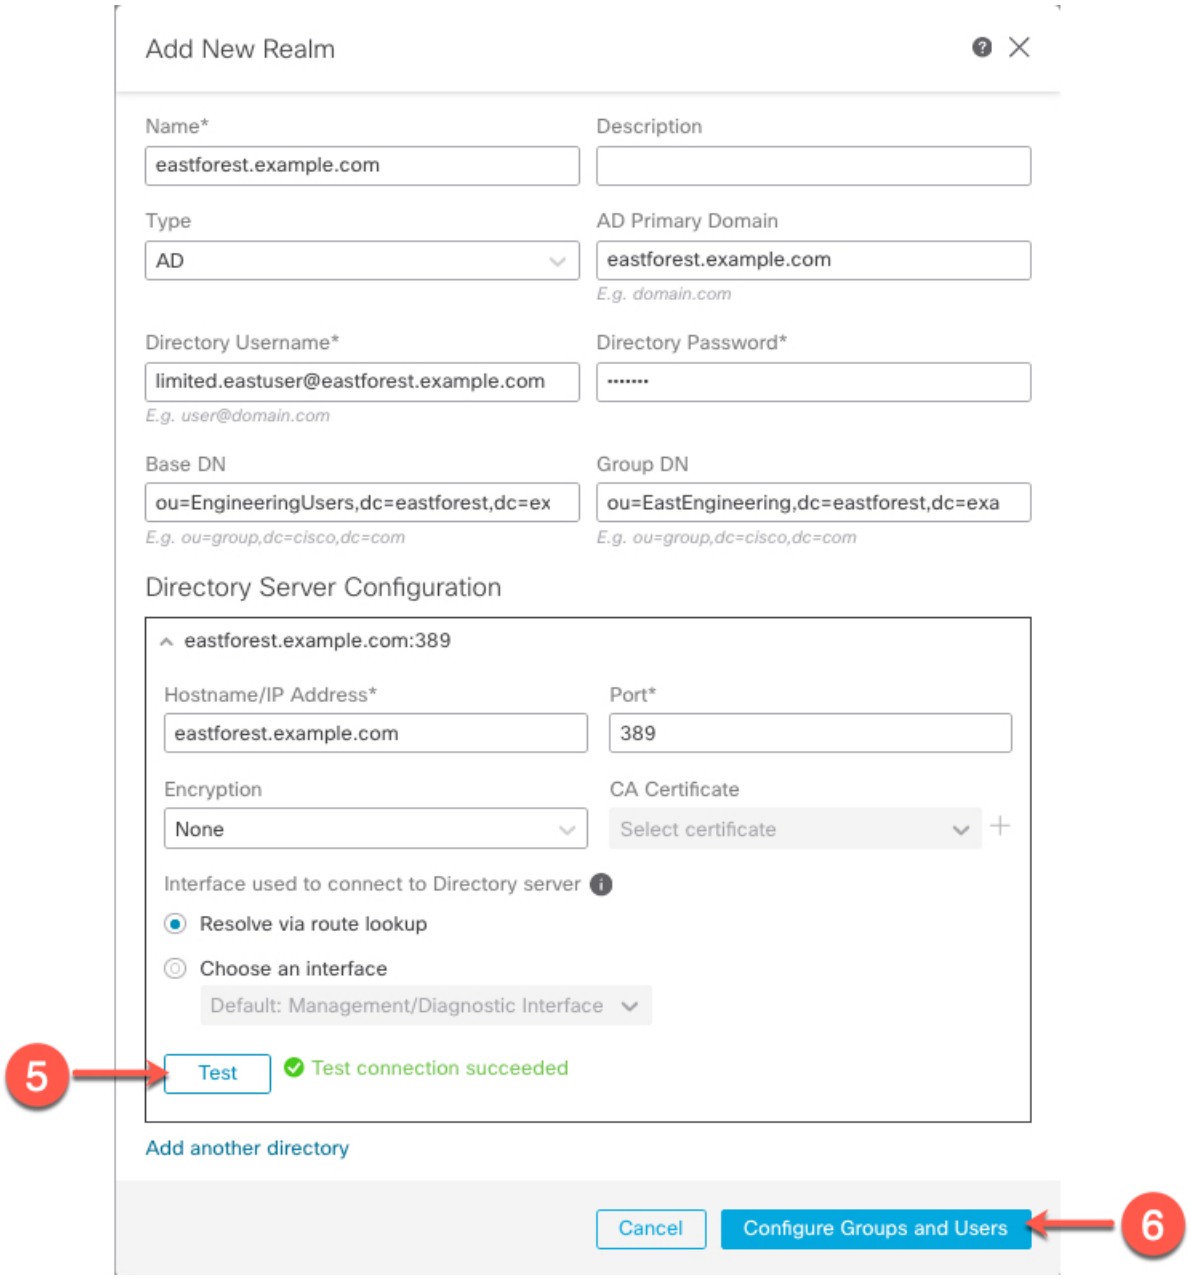

Note
The **Directory Username** can be any user in the Active Directory domain; no special permissions are required.

The **Interface used to connect to Directory server** can be any interface that can connect to the Active Directory server.

Step 5 Click **Test** and make sure the test succeeds before you continue.

Step 6 Click **Configure Groups and Users**.

Step 7 If your configuration was successful, the next page is displayed similar to the following.

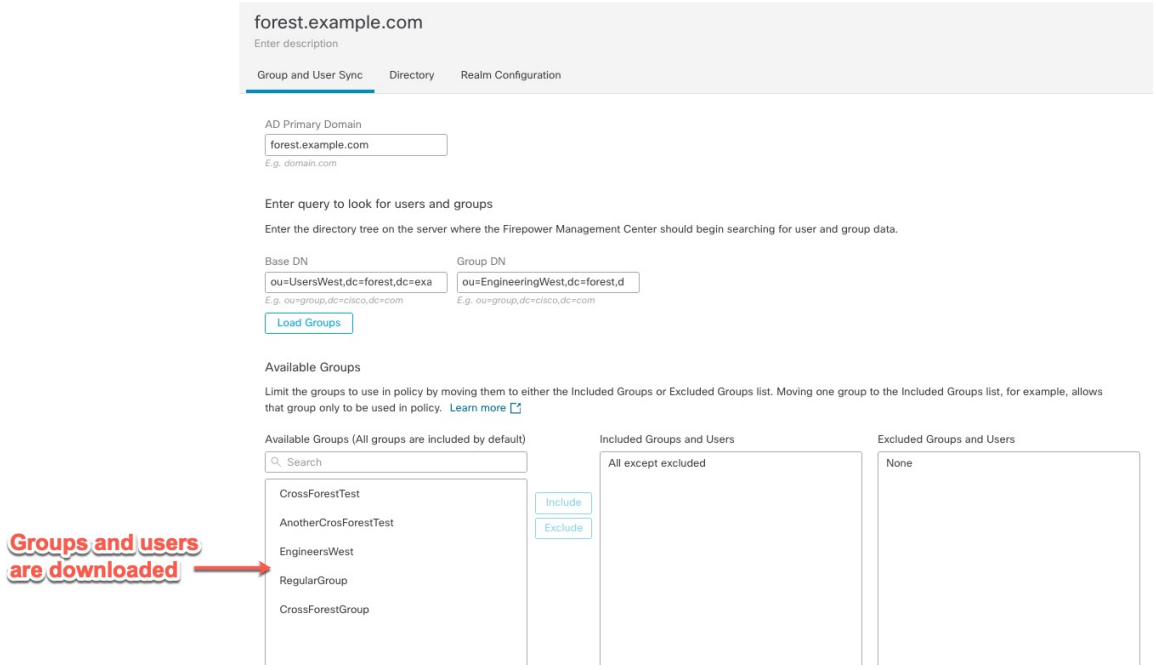

Note
If groups and users were not downloaded, verify the values in the **Base DN** and **Groups DN** fields and click **Load Groups**.

There are other optional configurations available on this page; for more information about them, see Realm Fields, on page 2436 and Realm Directory and Synchronize fields, on page 2440.

Step 8 If you made changes on this page or tab pages, click **Save**.

Step 9 Click **Integration** > **Other Integrations** > **Realms** > **Realms**.

Step 10 Click **Add Realm**.

Step 11 Enter the following information to configure **eastforest.example.com**.

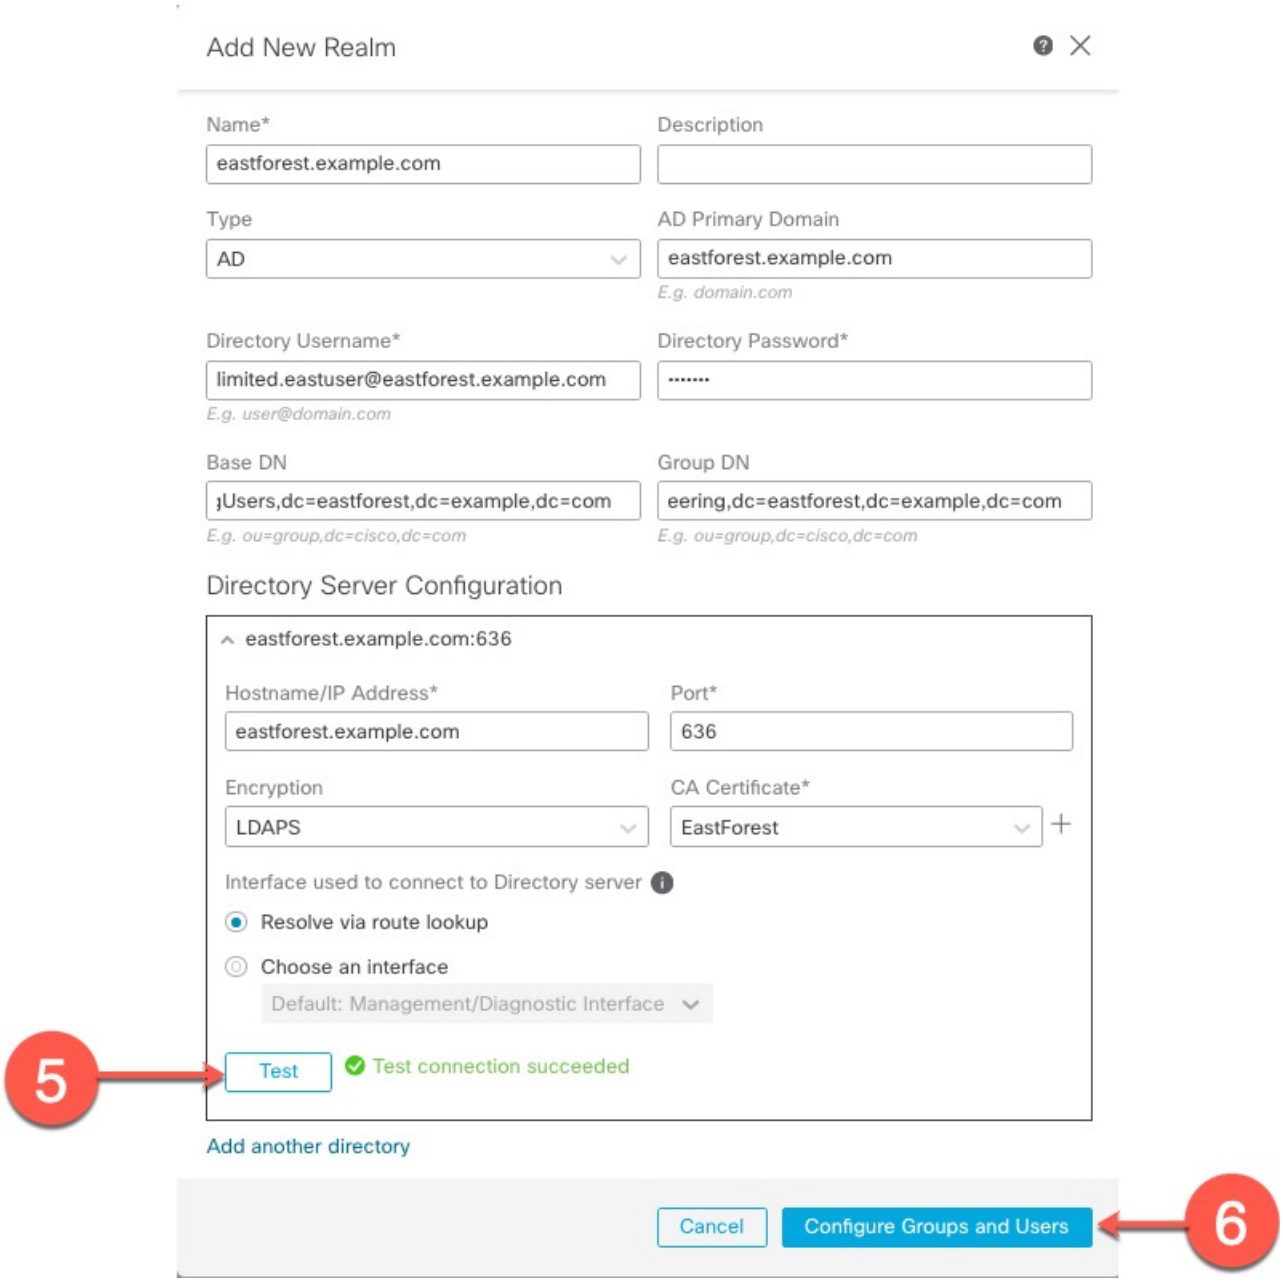

Step 12 Click **Test** and make sure the test succeeds before you continue.

Step 13 Click **Configure Groups and Users**.

Step 14 If your configuration was successful, the next page is displayed similar to the following.

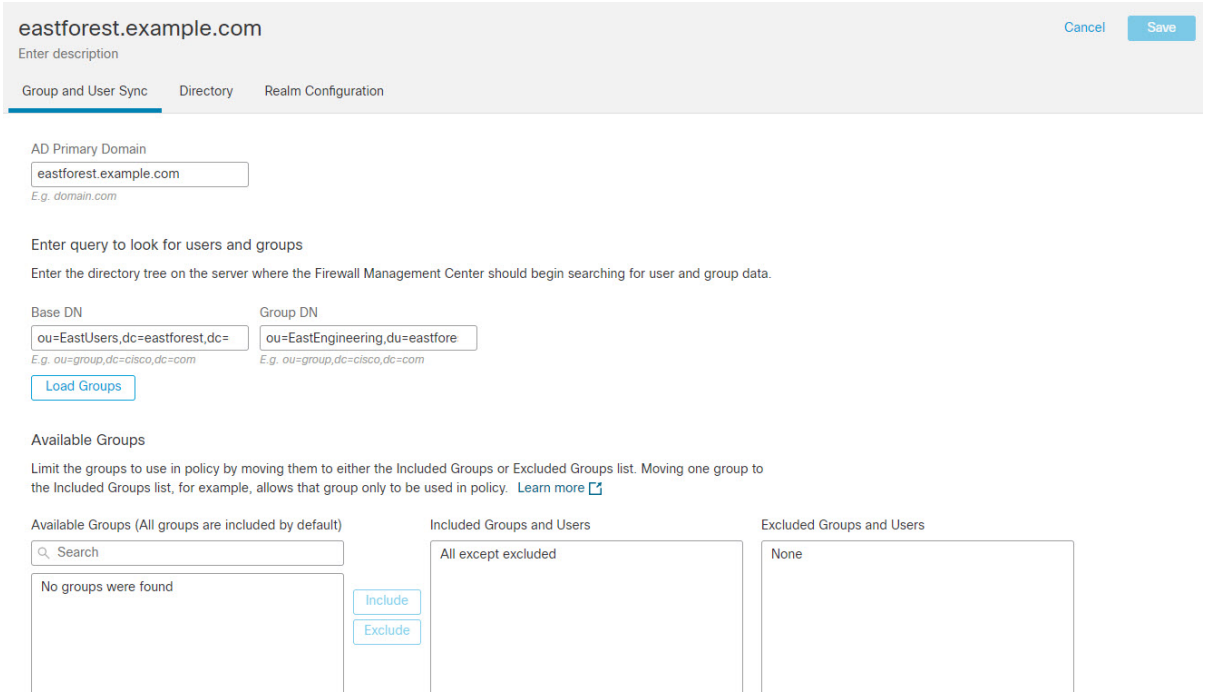

Related Topics

Configure the Secure Firewall Management Center for Cross-Domain-Trust Step 2: Synchronize Users and Groups, on page 2453

Configure the Secure Firewall Management Center for Cross-Domain-Trust Step 2: Synchronize Users and Groups

After you configure two or more Active Directory servers that have a cross-domain trust relationship, you must download users and groups. That process exposes possible issues with the Active Directory configuration (for example, groups or users downloaded for one Active Directory domain but not the other).

Before you begin

Make sure you have performed the tasks discussed in Configure the Secure Firewall Management Center for Cross-Domain-Trust Step 1: Configure Realms and Directories, on page 2448.

Procedure

Step 1 Log in to the Secure Firewall Management Center if you have not already done so.

Step 2 Click **Integration** > **Other Integrations** > **Realms** > **Realms**.

Step 3 At the end of the row of any realm in the cross-domain trust, click **Download Now**(), then click **Yes**.

Step 4 Click **Check Mark** () (Notifications) > **Tasks**.

If groups and users fail to download, try again. If subsequent attempts fail, review your realm and directory setup as discussed in Realm Fields, on page 2436 and Realm Directory and Synchronize fields, on page 2440.

Step 5 Click **Integration** > **Other Integrations** > **Realms** > **Sync Results**.

Related Topics

Configure the Secure Firewall Management Center for Cross-Domain-Trust Step 3: Resolve Issues, on page 2454

Configure the Secure Firewall Management Center for Cross-Domain-Trust Step 3: Resolve Issues

The final step in setting up cross-domain trust in the Firewall Management Center is to make sure users and groups are downloaded without errors. A typical reason why users and groups do not download properly is that the realms to which they belong have not been downloaded to the Firewall Management Center.

This topic discusses how to diagnose that a group referred in one forest to cannot be downloaded because the realm is not configured to find the group in the domain controller hierarchy.

Before you begin

Procedure

Step 1 Log in to the Secure Firewall Management Center if you have not already done so.

Step 2 Click **Integration** > **Other Integrations** > **Realms** > **Realms**.

In the Realms column, if **Yellow Triangle** (⚠) is displayed next to the name of a realm, you have issues that must be resolved. If not, your results are configured properly and you can quit.

Step 3 Download users and groups again from the realms that display issues.

a) Click the **Realms** tab.

b) Click ⤓ (Download Now), then click **Yes**.

Step 4 Click the **Sync Results** tab page.

If the **Yellow Triangle** (⚠) is displayed in the Realms column, click **Yellow Triangle** (⚠) next to the realm that has issues.

Step 5 In the middle column, click either **Groups** or **Users** to find more information.

Step 6 In the Groups or Users tab page, click **Yellow Triangle** (⚠) to display more information.

The right column should display enough information you can isolate the source of the issue.

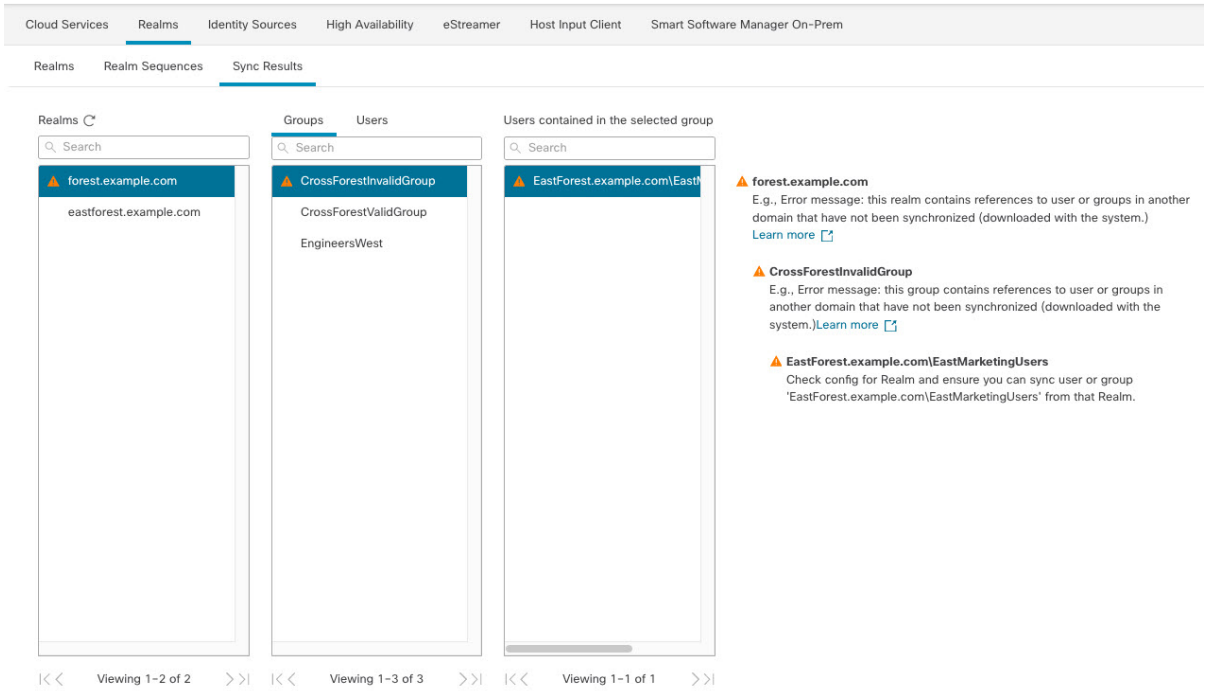

In the preceding example, **forest.example.com** includes a cross-domain group **CrossForestInvalidGroup** that contains another group **EastMarketingUsers** that was not downloaded by the Firewall Management Center. If, after synchronizing the **eastforest.example.com** realm again, the error does not resolve, it likely means that the Active Directory domain controller does not include **EastMarketingUsers**.

To resolve this issue, you can:

- Remove the **EastMarketingUsers** from **CrossForestInvalidGroup**, synchronize the **forest.example.com** realm again, and recheck.
- Remove the **ou=EastEngineering** value from the **Group DN** of the **eastforest.example.com** realm, which causes the Firewall Management Center to retrieve groups from the highest level in the Active Directory hierarchy, synchronize **eastforest.example.com**, and recheck.

Manage a Realm

This section discusses how to perform various maintenance tasks for a realm using controls on the Realms page:

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Procedure

Step 1 Log in to the Secure Firewall Management Center if you have not already done so.

Step 2 Click **Integration** > **Other Integrations** > **Realms** > **Realms**.

Step 3 To delete a realm, click **Delete** ( ).

Step 4 To edit a realm, click **Edit** ( ) next to the realm and make changes as described in Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.

Step 5 To enable a realm, slide **State** to the right; to disable a realm, slide it to the left.

Step 6 To download users and user groups, click **Download** ( ).

Step 7 To copy a realm, click **Copy** ( ).

Step 8 To compare realms, see Compare Realms, on page 2456.

Compare Realms

You must be an Admin, Access Admin, Network Admin, or Security Approver to perform this task.

Procedure

Step 1 Log in to the Secure Firewall Management Center if you have not already done so.

Step 2 Click **Integration** > **Other Integrations** > **Realms** > **Realms**.

Step 3 Click **Compare Realms**.

Step 4 Choose **Compare Realm** from the **Compare Against** list.

Step 5 Choose the realms you want to compare from the **Realm A** and **Realm B** lists.

Step 6 Click **OK**.

Step 7 To navigate individually through changes, click **Previous** or **Next** above the title bar.

Step 8 (Optional.) Click **Comparison Report** to generate the realm comparison report.

Step 9 (Optional.) Click **New Comparison** to generate a new realm comparison view.

Troubleshoot Realms and User Downloads

If you notice unexpected server connection behavior, consider tuning your realm configuration, device settings, or server settings. For other related troubleshooting information, see:

- Troubleshoot the ISE/ISE-PIC or Cisco TrustSec Issues, on page 2521
- Troubleshoot the TS Agent Identity Source, on page 2550
- Troubleshoot the Captive Portal Identity Source, on page 2541

- Troubleshoot the Remote Access VPN Identity Source, on page 2546
- Troubleshoot User Control, on page 2573

Symptom: Realms and groups reported but not downloaded

The Secure Firewall Management Center's health monitor informs you of user or realm mismatches, which are defined as:

- User mismatch: A user is reported to the Secure Firewall Management Center without being downloaded.

  A typical reason for a user mismatch is that the user belongs to a group you have excluded from being downloaded to the Secure Firewall Management Center. Review the information discussed in Cisco Secure Firewall Management Center Device Configuration Guide.

- Realm mismatch: A user logs into a domain that corresponds to a realm not known to the Firewall Management Center.

For example, if you defined a realm that corresponds to a domain named **domain.example.com** in the Firewall Management Center but a login is reported from a domain named **another-domain.example.com**, this is a *realm mismatch*. Users in this domain are identified by the Firewall Management Center as Unknown.

You set the mismatch threshold as a percentage, above which a health warning is triggered. Examples:

- If you use the default mismatch threshold of 50%, and there are two mismatched realms in eight incoming sessions, the mismatch percentage is 25% and no warning is triggered.
- If you set the mismatch threshold to 30% and there are three mismatched realms in five incoming sessions, the mismatch percentage is 60% and a warning is triggered.

Unknown users that do not match identity rules have no policies applied to them. (Although you can set up identity rules for Unknown users, we recommend keeping the number of rules to a minimum by identifying users and realms correctly.)

For more information, see Detect Realm or User Mismatches, on page 2460.

Symptom: Users are not downloaded

Possible causes follow:

- If you have the realm **Type** configured incorrectly, users and groups cannot be downloaded because of a mismatch between the attribute the system expects and what the repository provides. For example, if you configure **Type** as **LDAP** for a Microsoft Active Directory realm, the system expects the `uid` attribute, which is set to `none` on Active Directory. (Active Directory repositories use `sAMAccountName` for the user ID.)

  Solution: Set the realm **Type** field appropriately: **AD** for Microsoft Active Directory or **LDAP** for another supported LDAP repository.

- Users in Active Directory groups that have special characters in the group or organizational unit name might not be available for identity policy rules. For example, if a group or organizational unit name contains the characters asterisk (`*`), equals (`=`), or backslash (`\`), users in those groups are not downloaded and can't be used for identity policies.

  Solution: Remove special characters from the group or organizational unit name.

Important To reduce latency between the Secure Firewall Management Center and your Active Directory domain controller, we strongly recommend you configure a realm directory (that is, domain controller) that is as close as possible geographically to the Secure Firewall Management Center.

For example, if your Secure Firewall Management Center is in North America, configure a realm directory that is also in North America. Failure to do so can cause problems such as timeout downloading users and groups.

Symptom: Not all users in a realm are downloaded

Possible causes follow:

- If you attempt to download more than the maximum number of users in any one realm, the download stops at the maximum number of users and a health alert is displayed. User download limits are set per Secure Firewall Management Center model. For more information, see User Limits for Microsoft Active Directory, on page 2384.
- Every user must be a member of a group. Users that are members of no groups do not get downloaded.

Symptom: Access control policy doesn't match group membership

This solution applies to an AD domain that is in a trust relationship with other AD domains. In the following discussion, *external domain* means a domain other than the one to which the user logs in.

If a user belongs to a group defined in a trusted external domain, the Secure Firewall Management Center doesn't track membership in the external domain. For example, consider the following scenario:

- Domain controllers 1 and 2 trust each other
- Group A is defined on domain controller 2
- User `mparvinder` in controller 1 is a member of Group A

Even though user `mparvinder` is in Group A, the Secure Firewall Management Center access control policy rules specifying membership Group A don't match.

Solution: Create a similar group in domain controller 1 that contains has all domain 1 accounts that belong to group A. Change the access control policy rule to match any member of Group A or Group B.

Symptom: Access control policy doesn't match child domain membership

If a user belongs to a domain that is child of parent domain, Firepower doesn't track the parent/child relationships between domains. For example, consider the following scenario:

- Domain `child.parent.com` is child of domain `parent.com`
- User `mparvinder` is defined in `child.parent.com`

Even though user `mparvinder` is in a child domain, the Firepower access control policy matching the `parent.com` don't match `mparvinder` in the `child.parent.com` domain.

Solution: Change the access control policy rule to match membership in either `parent.com` or `child.parent.com`.

Symptom: Realm or realm directory test fails

The **Test** button on the directory page sends an LDAP query to the hostname or IP address you entered. If it fails, check the following:

- The **Hostname** you entered resolves to the IP address of an LDAP server or Active Directory domain controller.
- The **IP Address** you entered is valid.

The **Test AD Join** button on the realm configuration page verifies the following:

- DNS resolves the **AD Primary Domain** to an LDAP server or Active Directory domain controller's IP address.
- The **AD Join Username** and **AD Join Password** are correct.

  AD Join Username must be fully qualified (for example, **`administrator@mydomain.com`**, *not* **`administrator`**).
- The user has sufficient privileges to create a computer in the domain and join the Secure Firewall Management Center to the domain as a Domain Computer.

Symptom: User timeouts are occurring at unexpected times

If you notice the system performing user timeouts at unexpected intervals, confirm that the time on your ISE/ISE-PIC server is synchronized with the time on the Secure Firewall Management Center. If the appliances are not synchronized, the system may perform user timeouts at unexpected intervals.

If you notice the system performing user timeouts at unexpected intervals, confirm that the time on your ISE/ISE-PIC, or TS Agent server is synchronized with the time on the Secure Firewall Management Center. If the appliances are not synchronized, the system may perform user timeouts at unexpected intervals.

Symptom: User data for previously-unseen ISE/ISE-PIC users is not displaying in the web interface

After the system detects activity from an ISE/ISE-PIC or TS Agent user whose data is not yet in the database, the system retrieves information about them from the server. In some cases, the system requires additional time to successfully retrieve this information from Microsoft Windows servers. Until the data retrieval succeeds, activity seen by the ISE/ISE-PIC, or TS Agent user is *not* displayed in the web interface.

Note that this can also prevent the system from handling the user's traffic using access control rules.

Symptom: User data in events is unexpected

If you notice user or user activity events contain unexpected IP addresses, check your realms. The system does not support configuring multiple realms with the same **AD Primary Domain** value.

Symptom: Users originating from terminal server logins are not uniquely identified by the system

If your deployment includes a terminal server and you have a realm configured for one or more servers connected to the terminal server, you must deploy the Cisco Terminal Services (TS) Agent to accurately report user logins in terminal server environments. When installed and configured, the TS Agent assigns unique ports to individual users so the system can uniquely identify those users in the web interface.

For more information about the TS Agent, see the *Cisco Terminal Services (TS) Agent Guide*.

Detect Realm or User Mismatches

This section discusses how to detect realm or user *mismatches*, which are defined as:

- User mismatch: A user is reported to the Secure Firewall Management Center without being downloaded.

  A typical reason for a user mismatch is that the user belongs to a group you have excluded from being downloaded to the Secure Firewall Management Center. Review the information discussed in Cisco Secure Firewall Management Center Device Configuration Guide.

- Realm mismatch: A user logs into a domain that corresponds to a realm not known to the Firewall Management Center.

For additional details, see Troubleshoot Realms and User Downloads, on page 2456.

Unknown users that do not match identity rules have no policies applied to them. (Although you can set up identity rules for Unknown users, we recommend keeping the number of rules to a minimum by identifying users and realms correctly.)

Procedure

Step 1 Enable detection of realm or user mismatches:

a) Log in to the Firewall Management Center if you have not already done so.
b) Click **System** (⚙) > **Health** > **Policy**.
c) Create a new health policy or edit an existing one.
d) On the Editing Policy page, set a **Policy Runtime Interval**.
This is the frequency at which all health monitor tasks run.
e) In the left pane, click **Realm**.
f) Enter the following information:

- **Enabled**: Click **On**
- **Warning Users match threshold %**: The percentage of either realm mismatches or user mismatches that triggers a warning in the Health Monitor. For more information, see Troubleshoot Realms and User Downloads, on page 2456.

g) At the bottom of the page, click **Save Policy & Exit**.
h) Apply the health policy to managed devices as discussed in *Applying Health Policies* in the Cisco Secure Firewall Management Center Administration Guide.

Step 2 View user and realm mismatches in any of the following ways:

- If the warning threshold is exceeded, click **Warning** > **Health** in the top navigation of the Firewall Management Center. This opens the Health Monitor.
- Click **System** (⚙) > **Health** > **Monitor**.

Step 3 On the Health Monitor page, in the Display column, expand **Realm: Domain** or **Realm: User** to view details about the mismatch.

Troubleshoot Cross-Domain Trust

Typical issues with troubleshooting the Firewall Management Center configuration for cross-domain trust include the following:

- Not adding realms or directories for all forests that have shared groups.
- Configure a realm to exclude users from being downloaded and those users are referenced in a group in a different realm.
- Certain temporary issues.

Understand the issues

If there are issues with the Firewall Management Center being able to synchronize users and groups with your Active Directory forests, the Sync Results tab page is displayed similar to the following.

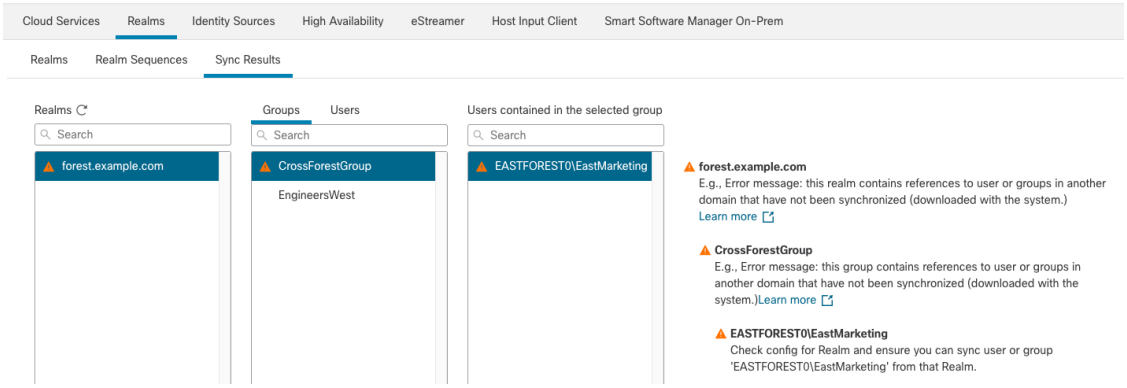

The following table explains how to interpret the information.

Column	Meaning
Realms	Displays all realms configured in the system. Click **Refresh** ( ) to update the list of realms. **Yellow Triangle** ( ) is displayed to indicate issues in the realm. Nothing is displayed next to a realm if all users and groups synchronized successfully.
Groups	Click **Groups** to display all groups in the realm. As with realms, **Yellow Triangle** ( ) is displayed to indicate issues. Click **Yellow Triangle** ( ) to see more detail about the issue.
Users	Click **Users** to display all users, sorted by group.
Users contained in the selected group	Displays all users in the group you selected in the Groups column. Clicking **Yellow Triangle** ( ) displays more information to the right of the table.

Column	Meaning
Groups that contain selected user	Displays all groups the selected user belongs to. Clicking **Yellow Triangle** ( ) displays more information to the right of the table.
Error detail information (displayed to the right of the table).	The system displays the NetBIOS forest name and group name it could not synchronize. Typical reasons the system cannot synchronize these users and groups follow: • **Problem**: The forest containing the groups and users do not have corresponding realms configured in the Firewall Management Center. **Solution**: Add a realm for the forest that contains the group as discussed in Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426. • **Problem**: You excluded groups from being downloaded to the Firewall Management Center. **Solution**: Click the **Realms** tab page, click **Edit** ( ), then move the indicated group or user from the **Excluded Groups and Users** list.

Try downloading users and groups again

If there is a possibility the issues are temporary, download users and groups for all realms.

1. If you haven't done so already, log in to the Firewall Management Center.
2. Click **Integration** > **Other Integrations** > **Realms** > **Realms**.
3. Click **Download** ( ).
4. Click the **Sync Results** tab page.
5. If no indicator is displayed for entries in the Realms column, the issues have been resolved.

Add a realm for all forests

Make sure you configured:

- Firewall Management Center realm for each forest that has users you want to use in identity policies.
- Firewall Management Center directory for each domain controller in that forest with users you want to use in identity polices.

The following figure shows an example.

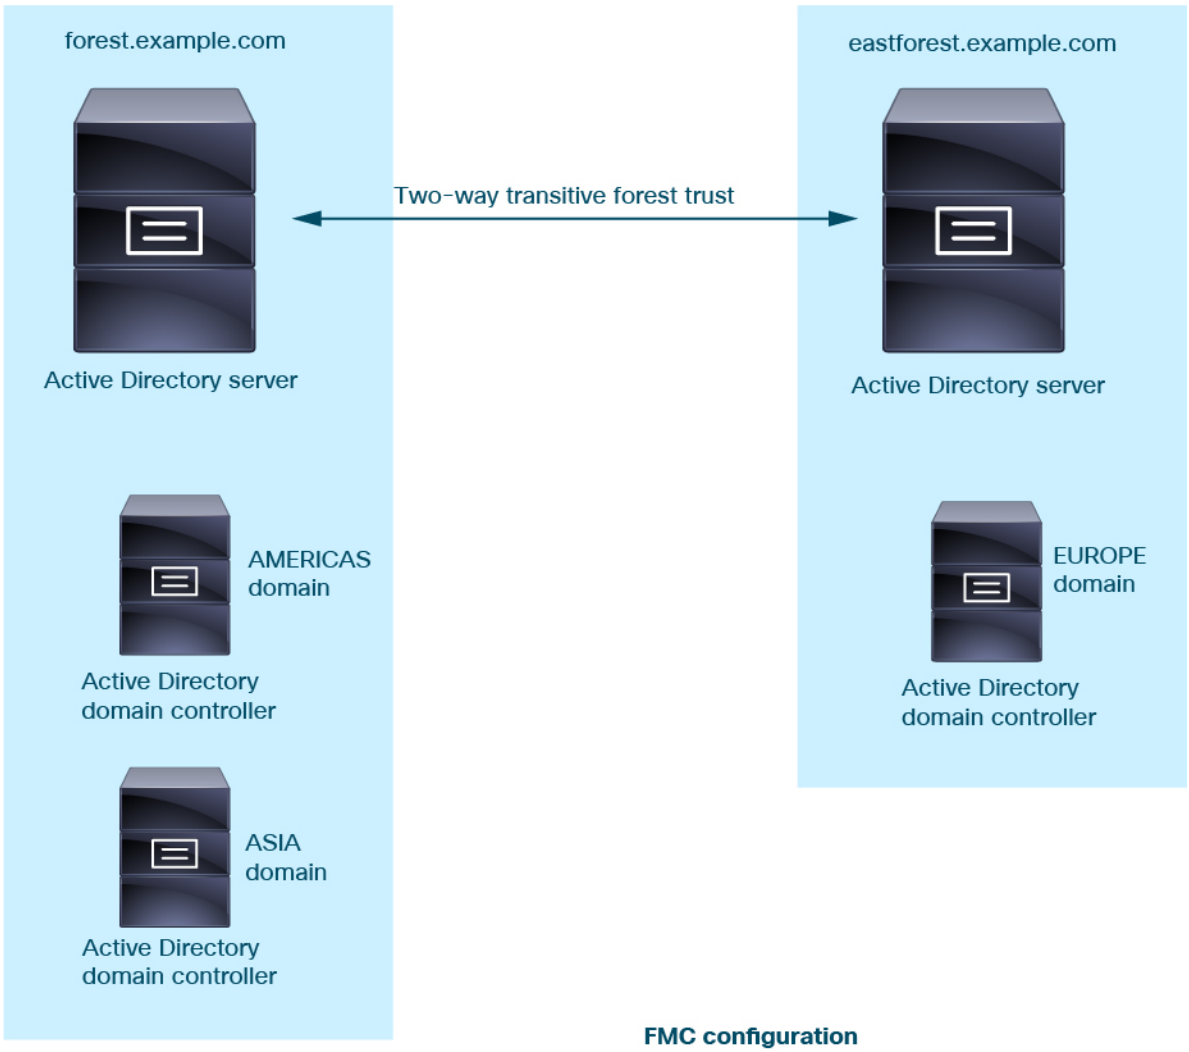

FMC configuration

Realm: forest.example.com
Directory: AMERICAS.forest.example.com
Directory: ASIA.forest.example.com

Realm: eastforest.example.com
Directory: EUROPE.eastforest.example.com

History for Realms

Table 136:

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Microsoft Azure Active Directory (SAML) realms.	7.6.0	7.4.0	You can now use Microsoft Azure Active Directory (AD) realms for active and passive authentication: • Active authentication using Azure AD: Use Azure AD as a captive portal. • Passive authentication using Cisco ISE (introduced in Version 7.4.0): The Firewall Management Center gets groups from Azure AD and logged-in user session data from ISE. We use SAML (Security Assertion Markup Language) to establish a trust relationship between a service provider (the devices that handle authentication requests) and an identity provider (Azure AD). **Upgrade impact**. If you had a Microsoft Azure AD realm configured before the upgrade, it is displayed as a SAML - Azure AD realm configured for passive authentication. All previous user session data is preserved. New/modified screens: **Integration** > **Other Integrations** > **Realms** > **Add Realm** > **SAML - Azure AD** New/modified CLI commands: none
Microsoft Azure Active Directory (AD) realms.	7.4.0	7.4.0	You can use a Microsoft Azure Active Directory (AD) realm with ISE to authenticate users and get user sessions for user control. New/modified screens: **System** > **Integration** > **Realms** > **Add Realm** > **Azure AD**
Proxy sequences.	7.2.0	7.2.0	Similar to a realm sequence, a proxy sequence is one or more managed devices that can communicate with Cisco Security Cloud Control in the event Cisco Security Cloud Control cannot communicate with the LDAP or Active Directory server. New/modified screens: **Integration** > **Other Integrations** > **Realms** > **Proxy Sequence**
Cross-domain trust for Active Directory domains.	7.2.0	7.0.0	A grouping of Microsoft Active Directory (AD) domains that trust each other is commonly referred to as a *forest*. This trust relationship can enable domains to access each other's resources in different ways. For example, a user account defined in domain A can be marked as a member of a group defined in domain B. The Firewall Management Center can get users from Active Directory forests for identity rules.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Realm sequences.	7.2.0	6.7.0	A *realm sequence* is an ordered list of two or more realms to which to apply identity rules. When you associate a realm sequence with an identity policy, the Firepower System searches the Active Directory domains in order from first to last as specified in the realm sequence. New/modified screens: **Integration** > **Other Integrations** > **Realms** > **Realm SequencesSystem** > **Integration** > **Realms** > **Realm Sequences**
Realms for user control.	7.2.0	Any	A realm is a connection between the Firewall Management Center either an Active Directory or LDAP user repository.

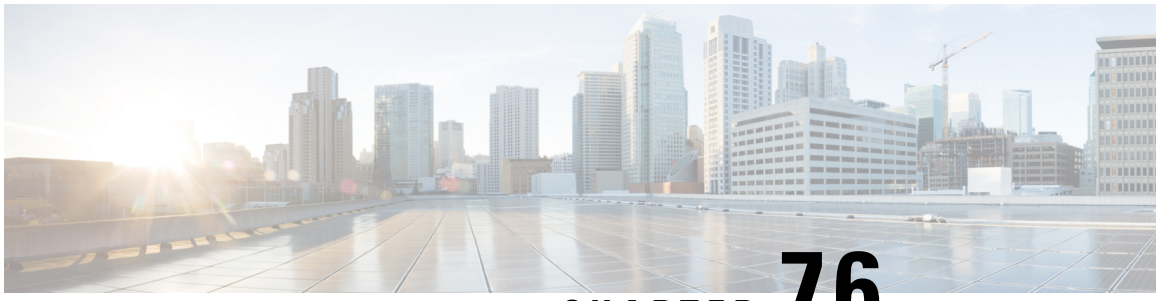

CHAPTER 76

User Control with the Passive Identity Agent

The following topics discuss how to configure and use the passive identity agent.

- The Passive Identity Agent Identity Source, on page 2467
- Deploy the Passive Identity Agent, on page 2469
- How to Create a Passive Identity Agent Identity Source, on page 2474
- Configure the Passive Identity Agent , on page 2476
- Monitor the Passive Identity Agent, on page 2491
- Manage the Passive Identity Agent, on page 2493
- Troubleshoot the Passive Identity Agent, on page 2494
- Security Requirements for the Passive Identity Agent, on page 2495
- Internet Access Requirements for the Passive Identity Agent, on page 2496
- History for the Passive Identity Agent, on page 2497

The Passive Identity Agent Identity Source

The passive identity agent identity source sends session data from Microsoft Active Directory (AD) to the Secure Firewall Management Center. All you need is a supported Microsoft AD setup as discussed in About Realms and Realm Sequences, on page 2429.

The passive identity agent version 1.0 sends IPv4 user sessions only but Version 1.1 sends IPv4 and IPv6 user sessions.

Note You do not need to configure the Cisco Identity Services Engine (ISE) to use this identity source.

Passive identity agent roles

The passive identity agent supports the following roles:

- Standalone: A passive identity agent that is not part of a redundant pair. A standalone agent can download users and groups from multiple Active Directory servers and domain controllers, provided the software is installed on all of them.
- Primary: (Primary agent in a redundant pair.) Can be installed on a Microsoft AD domain controller, directory server, or any network client.

Handles all communication with the Secure Firewall Management Center unless it stops communicating, in which case communication is handled by secondary agents.

- Secondary: (Secondary, or backup, agent in a redundant pair.) Can be installed on a Microsoft AD domain controller, directory server, or any network client.

  Monitors the health of the primary agent and takes over if the primary agent stops communicating with the Secure Firewall Management Center.

Passive identity agent system requirements

The passive identity agent requires the following:

- If you install the passive identity agent on a Windows Active Directory server, the server must run Windows Server 2008 or later.
- If you install it on a Windows client attached to the domain, the client must run Windows 8 or later.
- The system clock on all systems must be synchronized. We strongly recommend using the same NTP servers on all of them. This means:
  - The Secure Firewall Management Center.

    For more information, see Time Synchronization, on page 1002.
  - All Windows Active Directory servers and domain controllers.
  - The machine on which the passive identity agent is installed.
- Secure Firewall Management Center must run 7.6 or later.
- Any Secure Firewall Threat Defense managed by the Secure Firewall Management Center must run 7.1 or later.
- You must enable Snort 3 on the Secure Firewall Threat Defense devices.

Passive identity agent limitations

The passive identity agent the following limitations:

- Up to 10 agents simultaneously
- One passive identity agent identity source can monitor up to 50 AD directories
- Up to 300,000 concurrent user sessions
- IPv6 addresses are *not* supported (passive identity agent 1.0)
- IPv6 addresses are supported (passive identity agent 1.1)

Deploy the passive identity agent

For information about deployment options, see Deploy the Passive Identity Agent, on page 2469.

Note We recommend you use the latest version of the passive identity agent. To see the available versions, go to software.cisco.com. To upgrade the passive identity agent, see Upgrade the Passive Identity Agent Software, on page 2491

Deploy the Passive Identity Agent

You can install the Passive Identity Agent software on any machine that is part of a Microsoft Active Directory (AD) domain you want to use for user awareness and control. In other words, you can install it on any of the following:

- The Microsoft Active Directory server
- A domain controller
- A client connected to the network that is neither the directory server nor a domain controller

Any particular passive identity agent can monitor one or several Active Directory domain controllers in the same domain.

The machine on which the passive identity agent must communicate with the Secure Firewall Management Center using the TLS/SSL protocol. For more information, see Internet Access Requirements for the Passive Identity Agent, on page 2496.

Types of agents

You can configure the following types of agents on the Microsoft AD directory server, domain controller, or on any client connected to the domain:

- Standalone agent: One agent that can monitor one or several Active Directory domain controllers in the same domain.
- Primary agent and secondary agent that can monitor one or several AD domain controllers in the same domain: To provide redundancy, you can install a primary and secondary agent on different machines. The primary is responsible for communicating with the Secure Firewall Management Center but if communication fails, the secondary agent takes over.

See one of the following topics for more information.

Related Topics

Simple Passive Identity Agent Deployment, on page 2469
Single Passive Identity Agent Monitoring Multiple Domain Controllers, on page 2470
Multiple Passive Identity Agents Monitoring Multiple Domain Controllers, on page 2471
Passive Identity Agent Primary/Secondary Agent Deployments, on page 2473

Simple Passive Identity Agent Deployment

The following diagram shows the simplest passive identity agent deployment.

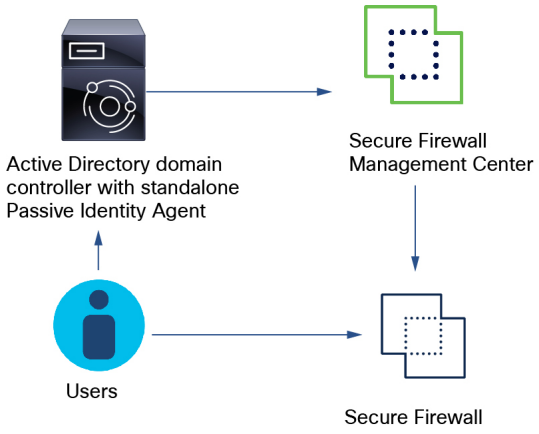

In the preceding example, a standalone passive identity agent is installed on the AD domain controller. Users log in and out of the AD domain and the agent sends user name and IP address information to the Secure Firewall Management Center. As users access the network, access control and identity policies deployed to the Secure Firewall Threat Defense determine whether or not, and how, access is allowed.

You can install a passive identity agent on the AD domain controller, directory server, or on any client connected to the domain you wish to monitor.

Single Passive Identity Agent Monitoring Multiple Domain Controllers

The following diagram shows a standalone passive identity agent that monitors several AD domain controllers.

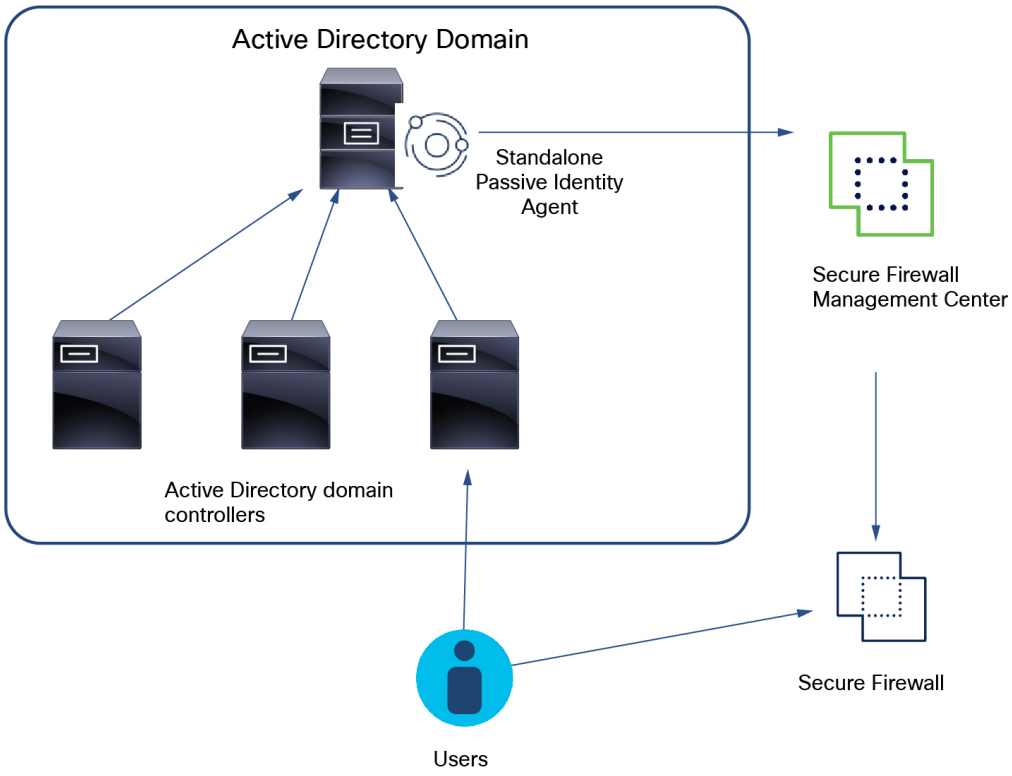

In the preceding diagram, the standalone passive identity agent is installed on a client attached to the AD domain (or on the domain controller itself). Users log in to any domain controller and the agent sends user and IP address information to the Secure Firewall Management Center. As users access the network, access control and identity policies deployed to the Secure Firewall Threat Defense determine whether or not, and how, access is allowed.

You can install a passive identity agent on the AD domain controller, directory server, or on any client connected to the domain you wish to monitor.

Multiple Passive Identity Agents Monitoring Multiple Domain Controllers

The following figure shows standalone monitoring multiple AD domain controllers:

- In AD domain 1, a standalone passive identity agent installed on a machine attached to AD domain controller 1 sends user and IP address mapping data to the Secure Firewall Management Center.
- In AD domain 2, standalone agents installed on AD domain controllers 1 and 2 send user and IP address mapping data to the Secure Firewall Management Center.

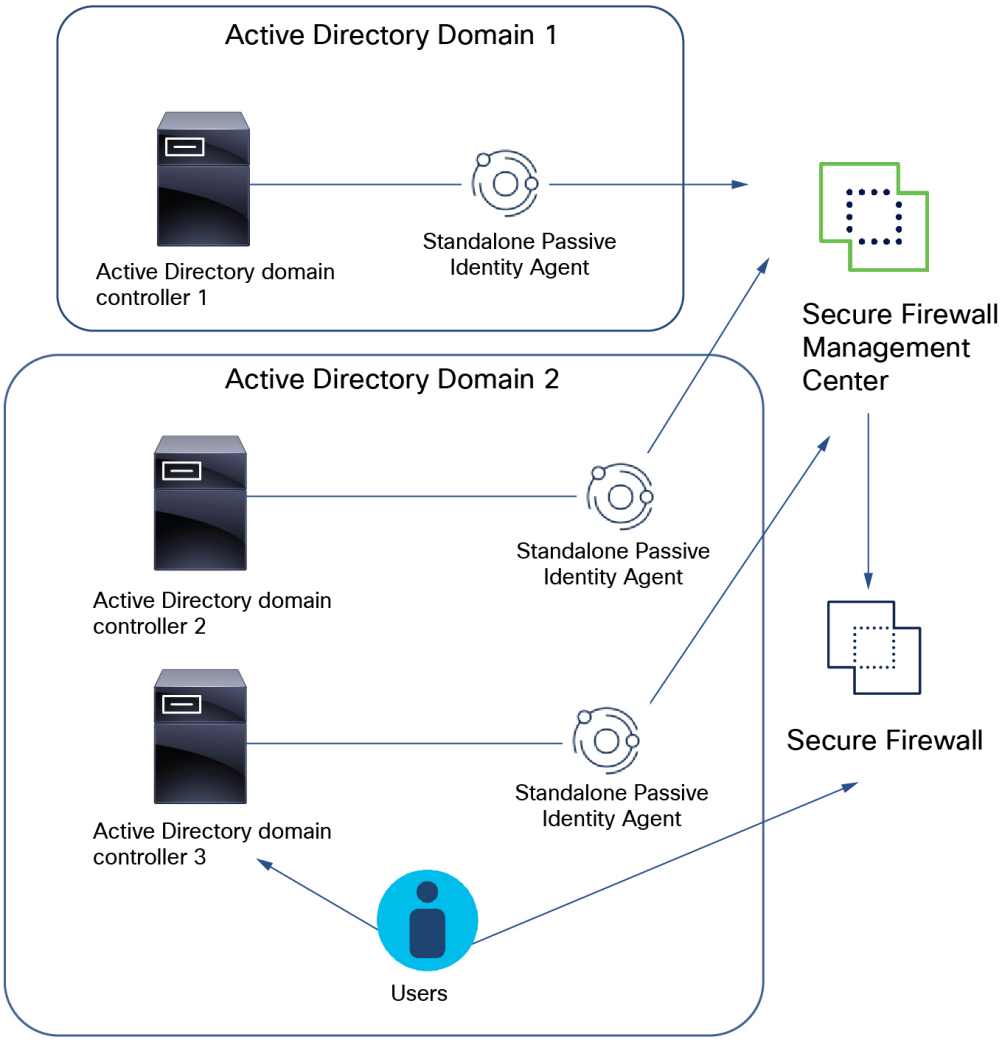

You can install a passive identity agent on the AD domain controller, directory server, or on any client connected to the domain you wish to monitor.

The preceding figure shows three passive identity agents, each configured as a standalone. To do this:

1. Create two Microsoft AD realms: one for each AD domain.

   See Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.

2. For AD domain 2, create two directories, one for each domain controller.

3. Install the Passive Identity Agent software on a client that can log in to the domain.

   Configure each passive identity agent individually to communicate with the Secure Firewall Management Center on which you configure the passive identity agent source.

   See Install the Passive Identity Agent Software, on page 2487.

4. Create the passive identity agent identity source.

   See Create a Primary or Secondary Passive Identity Agent Identity Source, on page 2479.

Passive Identity Agent Primary/Secondary Agent Deployments

To provide redundancy and to avoid a single point of failure, you can configure primary and secondary passive identity agents in any of the ways shown in this topic.

You can install a passive identity agent on the AD domain controller, directory server, or on any client connected to the domain you wish to monitor.

Single AD domain controller with primary and secondary agents

The following figure shows how to set up primary and secondary passive identity agents on one AD domain controller. If the primary agent fails, the secondary takes over.

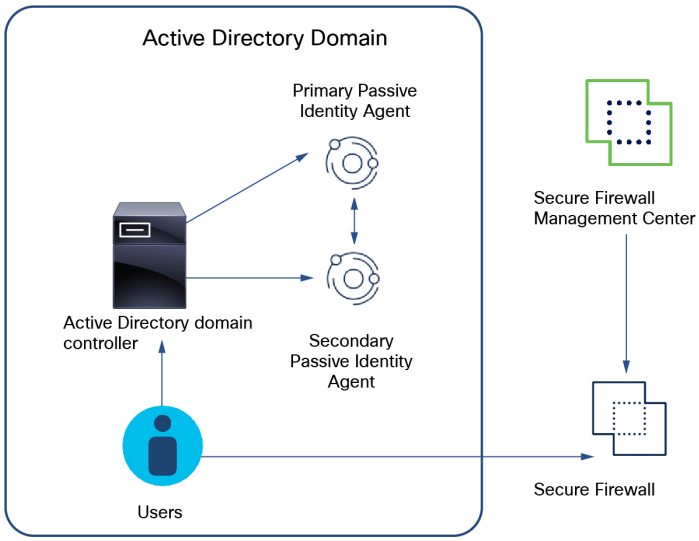

To set this up:

1. Create a Microsoft AD realm that has one directory for the domain controller.

   See Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.

2. Install the passive identity agent software on any two network machines connected to the domain controller.

   Configure each passive identity agent individually to communicate with the Secure Firewall Management Center on which you configure the passive identity agent source.

   See Install the Passive Identity Agent Software, on page 2487.

3. Create the identity source.

   See Create a Primary or Secondary Passive Identity Agent Identity Source, on page 2479.

Multiple AD domain controllers, primary and secondary agents

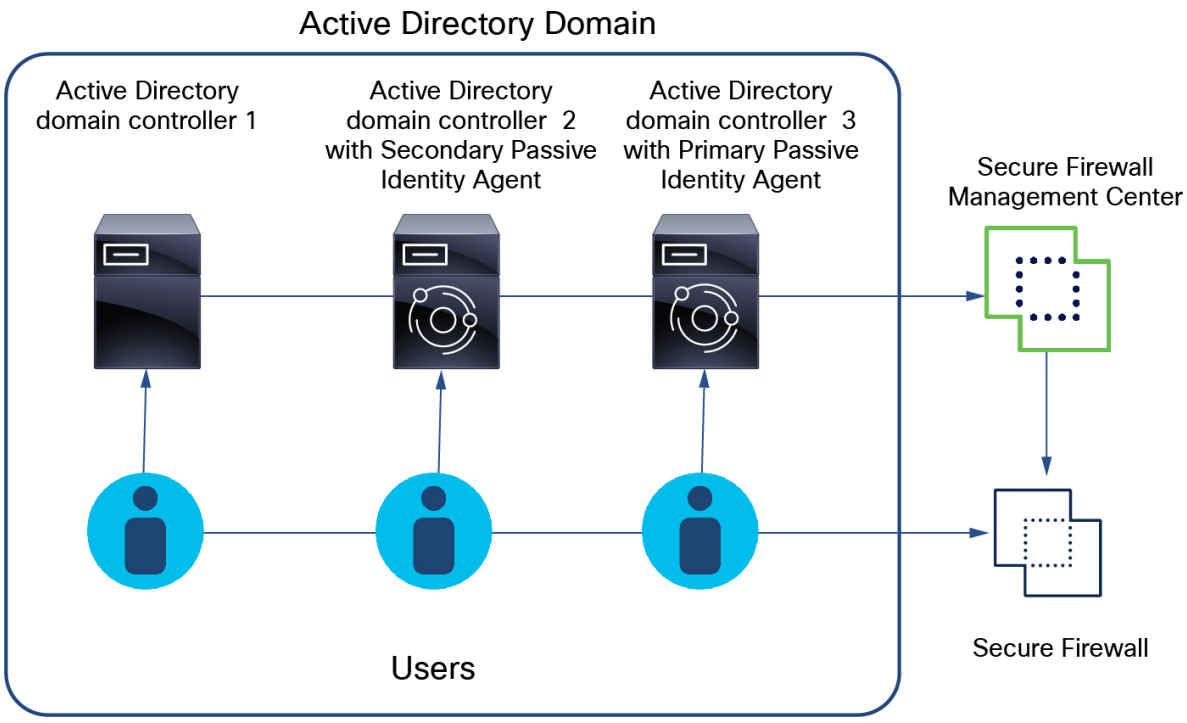

The preceding figure shows how to configure primary and secondary agents to monitor three AD domain controllers. If the primary agent fails, the secondary agent takes over.

To set this up:

1. Create a Microsoft AD realm that has one directory for the domain controller.

   See Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.

2. Install the passive identity agent software on any machine connected to the domain controller.

   Configure each passive identity agent individually to communicate with the Secure Firewall Management Center on which you configure the passive identity agent source.

   See Install the Passive Identity Agent Software, on page 2487.

3. Create the identity source.

   See Create a Primary or Secondary Passive Identity Agent Identity Source, on page 2479.

How to Create a Passive Identity Agent Identity Source

The following provides high-level tasks required to configure the passive identity agent identity source in the Secure Firewall Management Center and to deploy agent software to your Microsoft Active Directory (AD) servers.

Procedure

	Command or Action	Purpose
Step 1	Create a realm for your Microsoft AD domain and domain controllers.	*Realms* are connections between the Secure Firewall Management Center and the user accounts on the servers you monitor. They specify the connection settings and authentication filter settings for the server. For more information, see Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.
Step 2	Create a passive identity agent identity source.	The identity source allows the Secure Firewall Management Center and passive identity agent to communicate with each other. Create standalone, primary, or secondary agents, depending on your needs. For more information, see: • About Passive Identity Agent Roles, on page 2481 • Create a Passive Identity Agent Identity Source, on page 2476
Step 3	Create a passive identity agent user on the Secure Firewall Management Center.	We provide a role sufficient for the agent and manager to communicate with each other. We recommend using that role and no other for the passive identity agent user.
Step 4	Install the passive identity agent software.	The way you install the agent depends on your deployment. You can install a passive identity agent on the AD domain controller, directory server, or on any client connected to the domain you wish to monitor. For more information, see: • Deploy the Passive Identity Agent, on page 2469 • Install the Passive Identity Agent Software, on page 2487

What to do next

Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.

Configure the Passive Identity Agent

The following topics discuss how to configure the passive identity agent.

Related Topics

Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426
Create a Passive Identity Agent Identity Source, on page 2476
Install the Passive Identity Agent Software, on page 2487
Create a Secure Firewall Management Center User for the Passive Identity Agent, on page 2482

Create a Microsoft Active Directory Realm

The passive identity agent requires you to create a Microsoft Active Directory (AD) realm and directories in the Secure Firewall Management Center as discussed in Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.

Create a Passive Identity Agent Identity Source

This task discusses how to create a passive identity agent that sends user session activity to the Secure Firewall Management Center.

Before you begin

Complete the following:

- Review passive identity agent roles as discussed in About Passive Identity Agent Roles, on page 2481.
- Create a Microsoft AD realm as discussed in Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.

Procedure

Step 1 Log in to the Secure Firewall Management Center as an administrator.

Step 2 Click **Integration** > **Other Integrations** > **Identity Sources**.

Step 3 **Integration** > **Other Integrations** > **Identity Sources**.

Step 4 Click **Passive Identity Agent**.

Step 5 If the Cisco Secure Dynamic Attributes Connector has not been enabled yet, you are prompted to do so.

For more information about enabling the dynamic attributes connector, see Enable the Cisco Secure Dynamic Attributes Connector, on page 1806.

Step 6 Click **Create Agent**.

Step 7 In the Configure Agent dialog box, enter the following information:

Item	Description
Name	Enter a unique name to identify this passive identity agent.

Item	Description
Description	Enter an optional description.
Role	Click one of the following: • **Primary**: The agent responsible for communicating with the Secure Firewall Management Center. Not available if you choose **Standalone**. • **Secondary**: Becomes the primary if the primary loses contact with the Secure Firewall Management Center. Not available if you choose **Standalone**. • **Standalone**: If there is only one passive identity agent. For more information about roles, see About Passive Identity Agent Roles, on page 2481.

Step 8 Continue with:

- Create a Standalone Passive Identity Agent Identity Source, on page 2477
- Create a Primary or Secondary Passive Identity Agent Identity Source, on page 2479

Create a Standalone Passive Identity Agent Identity Source

This task discusses how to configure a standalone passive identity agent.

Before you begin

Complete the tasks discussed in Create a Passive Identity Agent Identity Source, on page 2476.

Procedure

Step 1 In the Configure Agent dialog box, enter the following information:

Item	Description
Role	Click **Standalone**.
Domain Controller	From the list, select the check box next to each domain controller that has a passive identity agent you wish to use for identity management and user control. (Optional.) Click **Add** (+) to add a new one.

The following figure shows an example of a standalone passive identity agent identity source.

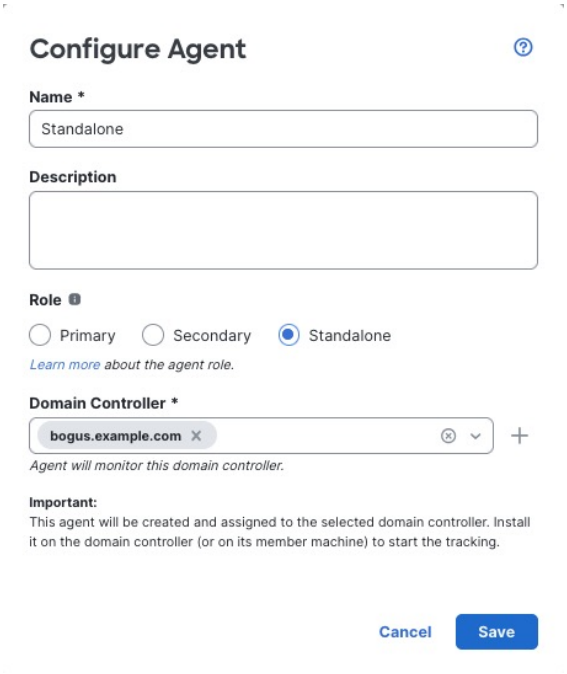

Step 2 In the Configure Agent dialog box, click **Save**.

Step 3 In the top right corner of the page, click **Save**.

The following figure shows an example.

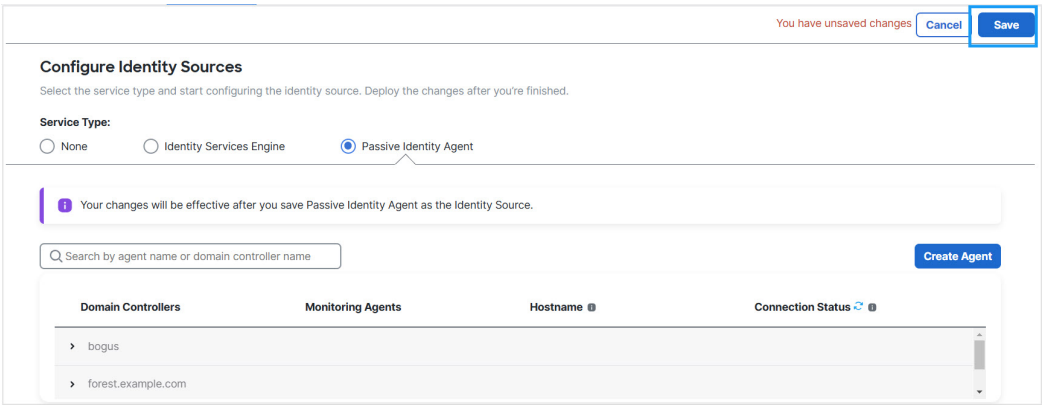

Note
The passive identity agent won't be active until you create a user and install the software.

What to do next

- See Create a Secure Firewall Management Center User for the Passive Identity Agent, on page 2482
- See About Passive Identity Agent Installation, on page 2485

Create a Primary or Secondary Passive Identity Agent Identity Source

The following task continues from Create a Passive Identity Agent Identity Source, on page 2476.

Before you begin

Complete the tasks discussed in Create a Passive Identity Agent Identity Source, on page 2476.

Procedure

Step 1 In the Configure Agent dialog box, enter the following information:

Item	Description
Role	Click one of the following: • **Primary**: The agent responsible for communicating with the Secure Firewall Management Center. • **Secondary**: Becomes the primary if the primary loses contact with the Secure Firewall Management Center. For more information about roles, see About Passive Identity Agent Roles, on page 2481.
Primary Agent Hostname/IP Address	(Primary agent only.) Enter the fully qualified domain name or IP address of the server on which the primary passive identity agent is installed. The passive identity agent version 1.0 supports IPv4 addresses and fully qualified domain names only. Version 1.1 supports IPv4, IPv6, and fully qualified domain names.
Secondary Agent Hostname/IP Address	(Secondary agent only.) Enter the fully qualified host name or IP address of the server on which the secondary passive identity agent is installed. The passive identity agent version 1.0 supports IPv4 addresses and fully qualified domain names only. Version 1.1 supports IPv4, IPv6, and fully qualified domain names.
Primary Agent	(Secondary agent only.) From the list, click the name of the primary passive identity agent.
Domain Controller	(Primary agent only.) From the list, select the check box next to each domain controller that has a passive identity agent you wish to use for identity management and user control.

The following figure shows an example of a primary agent:

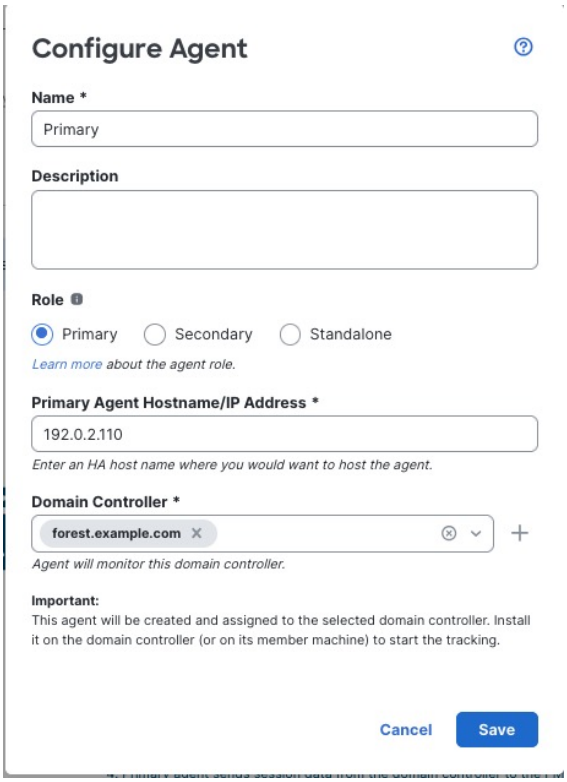

The following figure shows an example of a secondary agent:

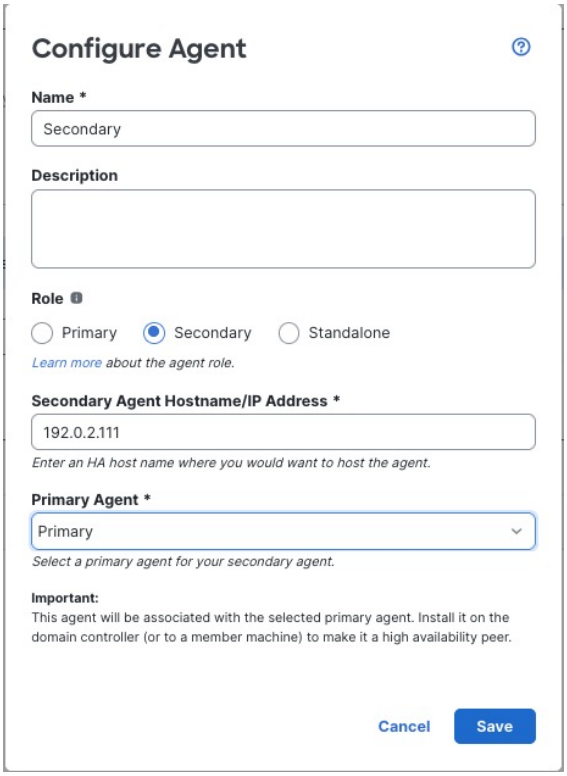

Step 2 In the Configure Agent dialog box, click **Save**.

Step 3 In the top right corner of the page, click **Save**.

The following figure shows an example.

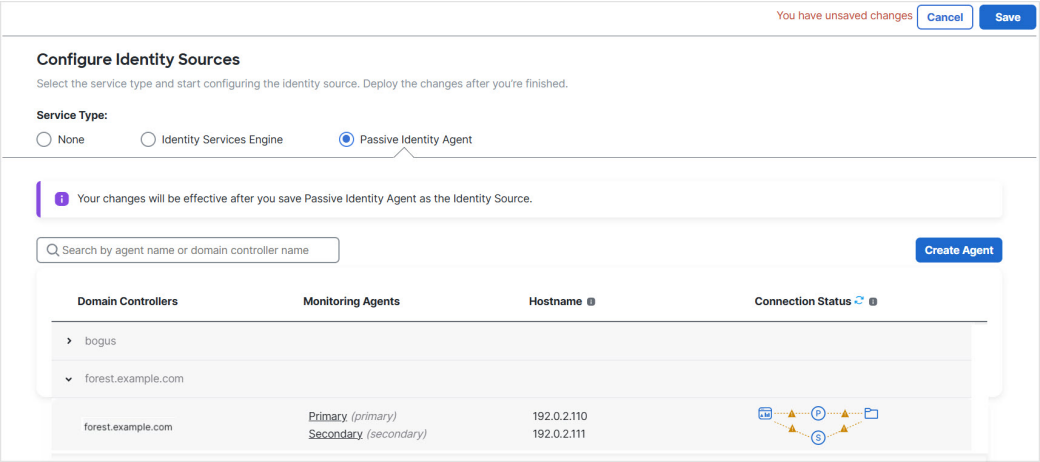

Note
The passive identity agent won't be active until you create a user and install the software.

What to do next

- See Create a Secure Firewall Management Center User for the Passive Identity Agent, on page 2482
- See About Passive Identity Agent Installation, on page 2485

About Passive Identity Agent Roles

The passive identity agent has the following roles:

- Standalone: A passive identity agent that is not part of a redundant pair. A standalone agent can download users and groups from multiple Active Directory servers and domain controllers, provided the software is installed on all of them.
- Primary: (Primary agent in a redundant pair.) Can be installed on a Microsoft AD domain controller, directory server, or any network client.

  Handles all communication with the Secure Firewall Management Center unless it stops communicating, in which case communication is handled by secondary agents.
- Secondary: (Secondary, or backup, agent in a redundant pair.) Can be installed on a Microsoft AD domain controller, directory server, or any network client.

  Monitors the health of the primary agent and takes over if the primary agent stops communicating with the Secure Firewall Management Center.

The can monitor several AD domain controllers that art part of the same domain.

Create a Secure Firewall Management Center User for the Passive Identity Agent

This task discusses how to create a Secure Firewall Management Center user with sufficient permissions to communicate with the passive identity agent. This user has limited privileges to perform other tasks; the user is expected only to enable communication with the passive identity agent.

Note Use *only* the **Passive Identity User** role for the passive identity agent user. In particular, *do not* use the **Administrator** role for the passive identity agent because **Administrator** will be logged off at a regular basis as the passive identity agent communicates with the Secure Firewall Management Center.

Before you begin

Complete the tasks discussed in Create a Passive Identity Agent Identity Source, on page 2476.

Note You *cannot* use external authentication with the Passive Identity Agent user.

Procedure

Step 1 Log in to the Secure Firewall Management Center as an administrator.

Step 2 Click **System** (⚙) > **Users** > **Users**.

Step 3 Click **Create User**.

Step 4 Create the user as discussed in Add or Edit an Internal User in the *Cisco Secure Firewall Management Center Administration Guide*.

Step 5 Select the **Passive Identity User** role.

The following figure shows an example.

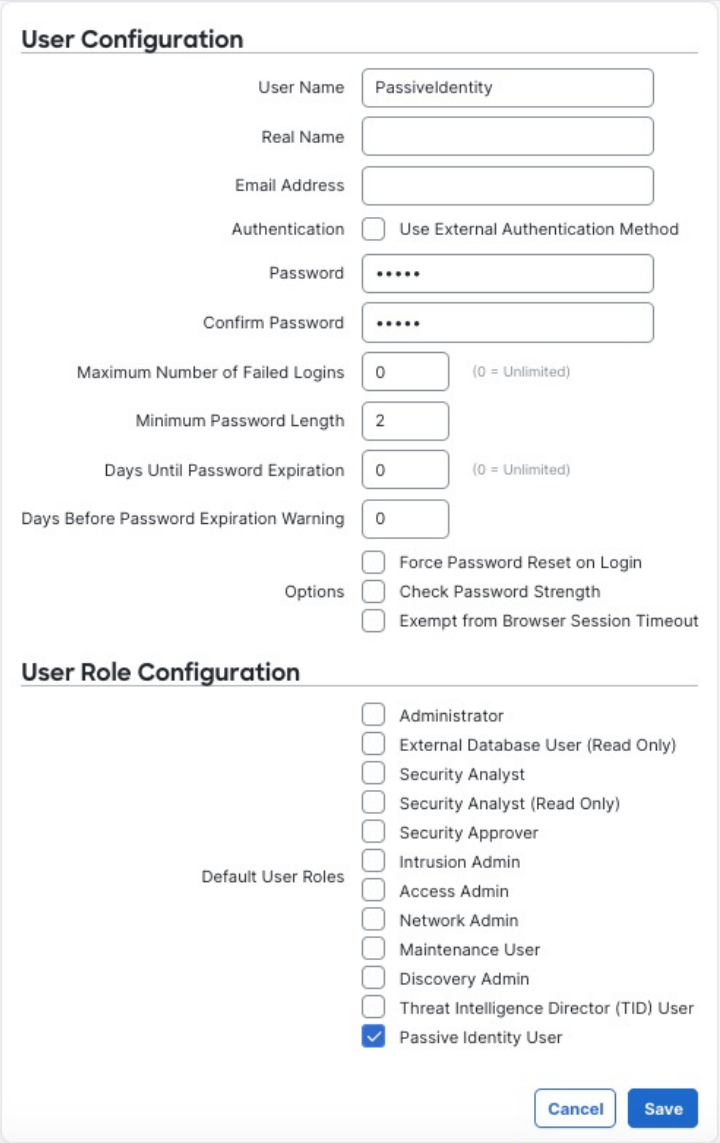

Note
Do not choose a role for the passive identity agent user other than **Passive Identity User** because the agent will not function properly.

Step 6 Click **Save**.

What to do next

About Passive Identity Agent Installation, on page 2485.

Troubleshoot the Passive Identity Agent

This topic discusses how you can troubleshoot the passive identity agent software on your Windows AD domain controller or directory server.

(Optional.) Set the log level

By default, the passive identity agent logs at the INFO level. To optionally change the log level, open **C:\Program Files\Program Files (x86)\Cisco\Cisco Passive Identity Agent\CiscoPassiveIdentityAgentService.exe.config** in a text editor, save the file, and restart the Cisco Passive Identity Agent service.

Do not rename the logging service

Do not rename **C:\Program Files\Program Files (x86)\Cisco\Cisco Passive Identity Agent\CiscoPassiveIdentityAgentService.exe.config**; otherwise, the passive identity agent will stop generating log files. Do not remove or change the **.exe.config** file extension.

Generate troubleshooting files

To generate a .zip containing troubleshooting files:

1. Log in to the Microsoft Active Directory domain controller.
2. Start the passive identity agent software.
3. Click the Troubleshooting button in the top right corner of the window.

   The following figure shows an example.

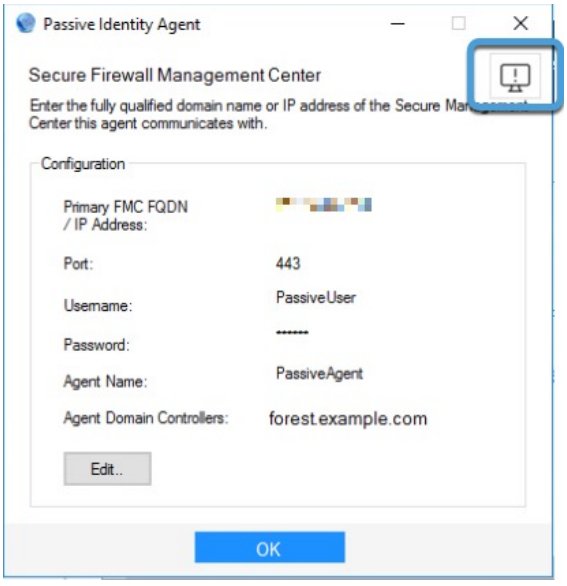

   A confirmation message is displayed.

Your troubleshoot logs are saved to your system's Downloads folder; the file name starts with **TroubleshootLogs**.

Manually view log files

Passive identity agent log files are stored in plain text format in the agent's installation directory: **C:\Program Files\Program Files (x86)\Cisco\Cisco Passive Identity Agent**.

Use Notepad or another text editor to view these files. Log files rotate after reaching 10MB in size.

Use the Microsoft Active Directory event viewer

In the event you are not seeing user sessions in the Secure Firewall Management Center, you can look on your Microsoft Active Directory server's event viewers for the following Kerberos-related events:

- 4770
- 4768

For general information about audit policy, see Audit Policy Recommendations on learn.microsoft.com.

For more information about Windows Group Policy Object settings, see Group Policy Objects on learn.microsoft.com.

About Passive Identity Agent Installation

The following topics discuss prerequisites and tasks required to install the passive identity agent.

Note We recommend you use the latest version of the passive identity agent. To see the available versions, go to software.cisco.com. To upgrade the passive identity agent, see Upgrade the Passive Identity Agent Software, on page 2491

Related Topics

Prerequisites to Installing the Passive Identity Agent, on page 2485
Install the Passive Identity Agent Software, on page 2487

Prerequisites to Installing the Passive Identity Agent

The following task shows how to configure Windows Group Policy Object (GPO) security settings to enable the passive identity agent to log successful and unsuccessful Kerberos authentication attempts. This setting is required for the passive identity agent to function properly.

For more information, see Audit Policy Recommendations on learn.microsoft.com.

Before you begin

Make sure your systems meet the following requirements:

- If you install the passive identity agent on a Windows Active Directory server, the server must run Windows Server 2008 or later.
- If you install it on a Windows client attached to the domain, the client must run Windows 8 or later.
- The system clock on all systems must be synchronized. We strongly recommend using the same NTP servers on all of them. This means:
  - The Secure Firewall Management Center.

For more information, see Time Synchronization, on page 1002.

- All Windows Active Directory servers and domain controllers.
- The machine on which the passive identity agent is installed.

- Secure Firewall Management Center must run 7.6 or later.
- Any Secure Firewall Threat Defense managed by the Secure Firewall Management Center must run 7.1 or later.
- You must enable Snort 3 on the Secure Firewall Threat Defense devices.

Procedure

Step 1 Log in to the Active Directory Server as an administrator.

Step 2 As Administrator, open a DOS command prompt.

Step 3 Enter `gpmc.msc` to start the Group Policy Management Editor.

Step 4 If necessary, create a new GPO; if one already exists, edit it.

For more information about creating a GPO, see a resource like Create a Group Policy Object on learn.microsoft.com.

Step 5 In your GPO, expand **Computer Configuration** > **Policies** > **Windows Settings** > **Security Settings** > **Advanced Policy Configuration** > **Audit Policies**.

Step 6 Click **Account Logon**.

Step 7 In the right pane, double-click **Audit Kerberos Authentication Service**.

Step 8 In the dialog box that is displayed, select all checkboxes which enables the system to log successes and failures.

The following figure shows an example.

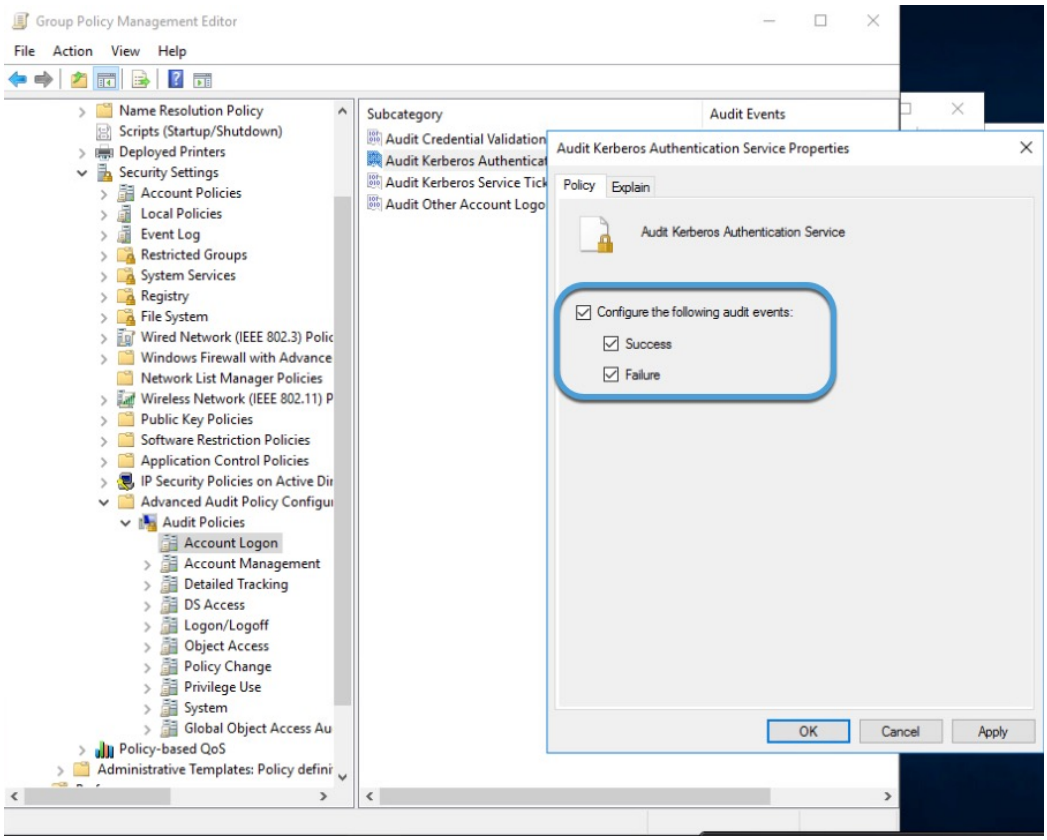

Step 9 Follow the prompts on your screen to save the changes.

Step 10 (Optional.) To update GPO immediately, enter **`gpupdate /force`** in your DOS command prompt window.

What to do next

See Install the Passive Identity Agent Software, on page 2487.

Install the Passive Identity Agent Software

This task discusses how to install the passive identity agent software. For a simple installation, you can install it on your Microsoft Active Directory (AD) domain controller; for other options, see Deploy the Passive Identity Agent, on page 2469.

Note We recommend you use the latest version of the passive identity agent. To see the available versions, go to software.cisco.com. To upgrade the passive identity agent, see Upgrade the Passive Identity Agent Software, on page 2491

Before you begin

See Prerequisites to Installing the Passive Identity Agent, on page 2485.

Procedure

Step 1 Download the passive identity agent from software.cisco.com.

Step 2 Log in as a member of the Administrators group to the machine on which to install the passive identity agent.

Step 3 Double-click **CiscoPassiveIdentityAgentInstaller-1.1.msi**.

Step 4 Click **Next**.

Step 5 Choose a folder in which to install the passive identity agent and click **Next**.

The default installation folder is **Program Files\Program Files (x86)\Cisco\Cisco Passive Identity Agent**.

Step 6 Click **Next**.

Step 7 Click **Install**.

Step 8 When the installation is done, click **Finish** and optionally check the box to start the passive identity agent.

Step 9 When the passive identity agent starts, click the **On-Prem** tab if you are using the agent with an on-premises Secure Firewall Management Center (physical or virtual) or click the **Cloud** tab if you are using the agent with Cloud-Delivered Firewall Management Center.

Step 10 In the Cisco Passive Agent dialog box, enter the following information:

Item	Description
FMC FQDN / IP Address	Enter the address of the Secure Firewall Management Center on which you created the passive identity agent identity source. The passive identity agent version 1.0 supports IPv4 addresses and fully qualified domain names only. Version 1.1 supports IPv4, IPv6, and fully qualified domain names.
Port	Enter the Secure Firewall Management Center listen port (by default, 443).
Username	Enter the username of the user you created in Create a Secure Firewall Management Center User for the Passive Identity Agent, on page 2482.
Password	Enter the user's password.
Agent	Click the list to locate the domain controller of the passive identity agent you created previously on the Secure Firewall Management Center.

The following figure shows an example.

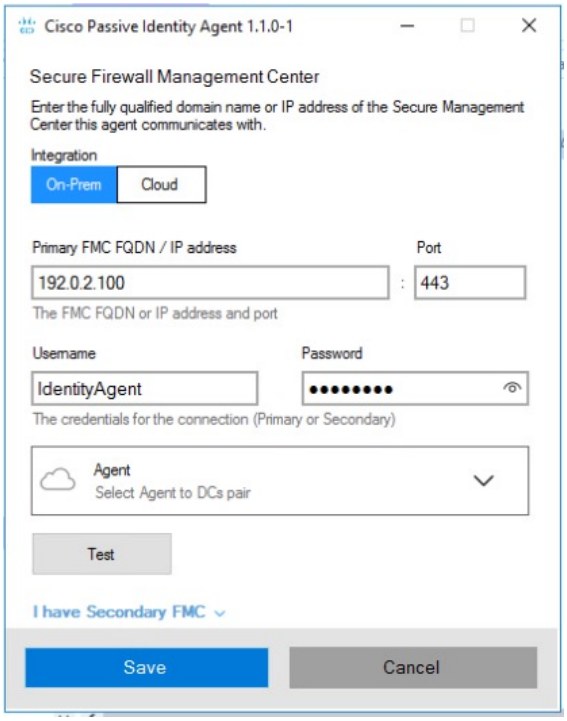

Step 11 Click the **Agent** list.

Step 12 From the list, click the name of the domain controller to monitor.

Step 13 Click **Test**.

The following figure shows an example.

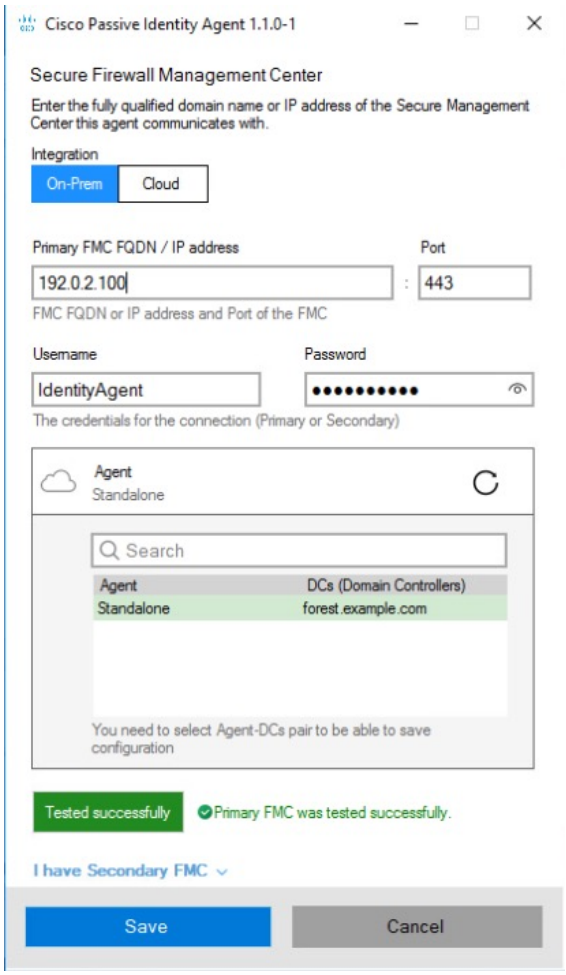

Step 14 If you have a high availability pair, click **I have Secondary FMC** and enter the secondary's IP address or fully qualified host name and its listen port.

Step 15 Only if the test succeeds, click **Save**.

What to do next

- Specify users to control and other options using an identity policy as described in Create an Identity Policy, on page 2555.
- Associate the identity rule with an access control policy, which filters and optionally inspects traffic, as discussed in Associating Other Policies with Access Control, on page 1762.
- Deploy your identity and access control policies to managed devices as discussed in Deploy Configuration Changes, on page 259.
- Monitor user activity as discussed in *Using Workflows* in the Cisco Secure Firewall Management Center Administration Guide.

Uninstall the Passive Identity Agent Software

This task discusses how to uninstall the passive identity agent software from your Microsoft AD servers.

Procedure

Step 1 Log in as an administrator to the machine on which the passive identity agent is installed.

Step 2 Search for Add or Remove Programs.

Step 3 Click **Cisco Passive Identity Agent**.

Step 4 Click **Uninstall**.

Step 5 You are required to confirm the uninstallation.

Upgrade the Passive Identity Agent Software

To upgrade to a newer version of the passive identity agent, you must uninstall the earlier version the install the later version.

See the following information:

- Uninstall the Passive Identity Agent Software, on page 2491
- About Passive Identity Agent Installation, on page 2485

Monitor the Passive Identity Agent

The passive identity agent indicates whether or not it can communicate with the Secure Firewall Management Center and other agents if it's configured as primary-secondary. You can view the status at **Integration** > **Other Integrations** > **Identity Sources**.

Deployments

A standalone passive identity agent is represented as follows.

A primary-secondary pair is represented as follows.

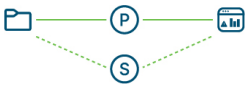

The following table explains the meaning of the indicators.

Object	Meaning
	Secure Firewall Management Center

Object	Meaning
	Standalone Passive Identity Agent
	Active Directory domain controller
	Primary agent
	Secondary agent

Status indicators and colors

The passive identity agent indicates status using lines (that indicate whether communication with the Secure Firewall Management Center is active or standby) and colors (that indicate whether or not communication is successful).

The following table shows the meanings of lines and colors:

Object	Meaning
Solid line	The agent that is responsible for communicating with the Secure Firewall Management Center.
Dashed line	Primary/secondary configuration only. The agent that is acting as the backup agent. In the event of a communication failure between the active (solid line) agent, this agent communicates with the Secure Firewall Management Center.
Blue	Agent communication is normal.
Amber	Agent has never successfully communicated with the Secure Firewall Management Center. A newly created agent line is amber and remains so until configuration is complete.
Red	Communication is failing. To resolve the issues: • Check sure the network connections between agents and the Secure Firewall Management Center. • Make sure you have completed configuring the system (Microsoft AD server, domain controllers, and the Secure Firewall Management Center). For more information, see How to Create a Passive Identity Agent Identity Source, on page 2474.

Manage the Passive Identity Agent

The following topics discuss how to edit or delete passive identity agents you previously configured on the Secure Firewall Management Center.

Related Topics

Edit Passive Identity Agents, on page 2493
Delete a Standalone Passive Identity Agent, on page 2493
Delete Primary and Secondary Passive Identity Agents, on page 2494
Uninstall the Passive Identity Agent Software, on page 2491

Edit Passive Identity Agents

This task discusses how to edit passive identity agents you previously configured in the Secure Firewall Management Center.

Procedure

Step 1 Log in to the Secure Firewall Management Center as an administrator.
Step 2 Click **Integration** > **Other Integrations** > **Identity Sources**.
Step 3 Click **Passive Identity Agent**.
Step 4 Click **Edit** (✎) next to the agent to edit.
Step 5 Make the desired changes.
Step 6 Click **Save**.

Delete a Standalone Passive Identity Agent

This task discusses how to delete a standalone passive identity agent.

Procedure

Step 1 Log in to the Secure Firewall Management Center as an administrator.
Step 2 Click **Integration** > **Other Integrations** > **Identity Sources**.
Step 3 Click **Passive Identity Agent**.
Step 4 Click **Edit** (✎) next to the agent to delete.
Step 5 Click **Delete**.
Step 6 You are required to confirm the action.

Delete Primary and Secondary Passive Identity Agents

This task discusses how to delete primary and secondary passive identity agents. You must delete a secondary agent before you can delete a primary agent.

Procedure

Step 1 Log in to the Secure Firewall Management Center as an administrator.

Step 2 Click **Integration** > **Other Integrations** > **Identity Sources**.

Step 3 Click **Passive Identity Agent**.

Step 4 Click **Edit** (✎) next to a secondary agent to delete.

Step 5 Click **Delete**.

Step 6 You are required to confirm the action.

Step 7 If you wish to delete a primary agent, first delete all secondary agents.

Troubleshoot the Passive Identity Agent

This topic discusses how you can troubleshoot the passive identity agent software on your Windows AD domain controller or directory server.

(Optional.) Set the log level

By default, the passive identity agent logs at the INFO level. To optionally change the log level, open **C:\Program Files\Program Files (x86)\Cisco\Cisco Passive Identity Agent\CiscoPassiveIdentityAgentService.exe.config** in a text editor, save the file, and restart the Cisco Passive Identity Agent service.

Do not rename the logging service

Do not rename **C:\Program Files\Program Files (x86)\Cisco\Cisco Passive Identity Agent\CiscoPassiveIdentityAgentService.exe.config**; otherwise, the passive identity agent will stop generating log files. Do not remove or change the **.exe.config** file extension.

Generate troubleshooting files

To generate a .zip containing troubleshooting files:

1. Log in to the Microsoft Active Directory domain controller.
2. Start the passive identity agent software.
3. Click the Troubleshooting button in the top right corner of the window.

   The following figure shows an example.

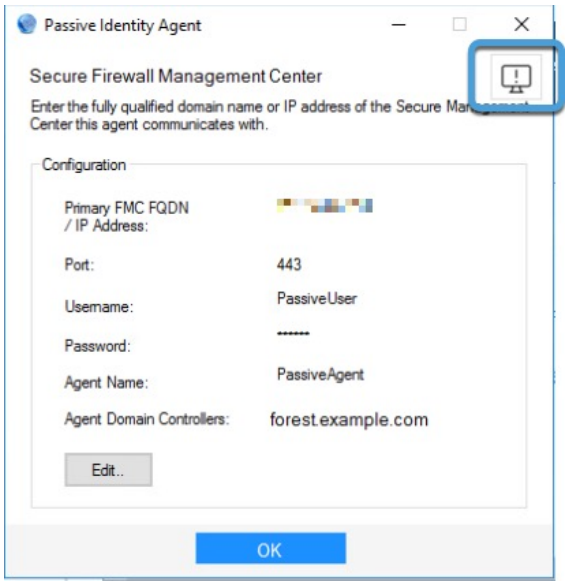

A confirmation message is displayed.

Your troubleshoot logs are saved to your system's Downloads folder; the file name starts with **TroubleshootLogs**.

Manually view log files

Passive identity agent log files are stored in plain text format in the agent's installation directory: **C:\Program Files\Program Files (x86)\Cisco\Cisco Passive Identity Agent**.

Use Notepad or another text editor to view these files. Log files rotate after reaching 10MB in size.

Use the Microsoft Active Directory event viewer

In the event you are not seeing user sessions in the Secure Firewall Management Center, you can look on your Microsoft Active Directory server's event viewers for the following Kerberos-related events:

- 4770
- 4768

For general information about audit policy, see Audit Policy Recommendations on learn.microsoft.com.

For more information about Windows Group Policy Object settings, see Group Policy Objects on learn.microsoft.com.

Security Requirements for the Passive Identity Agent

To safeguard the system, you should install the passive identity agent on a protected internal network. Although the passive identity agent is configured to have only the necessary services and ports available, you must make sure that attacks cannot reach it.

If the passive identity agent and the Secure Firewall Management Center reside on the same network, you can connect the Secure Firewall Management Center to the same protected internal network as the passive identity agent.

Regardless of how you deploy your appliances, inter-system communication is encrypted. However, you must still take steps to ensure that communications between appliances cannot be interrupted, blocked, or tampered with; for example, with a distributed denial of service (DDoS) or man-in-the-middle attack.

Internet Access Requirements for the Passive Identity Agent

By default, the passive identity agent is configured to communicate with the Firepower System over the internet using HTTPS on port 443/tcp (HTTPS). If you do not want the passive identity agent to have direct access to the internet, you can configure a proxy server.

The following information informs you of the ports the passive identity agent use to communicate with each other, with the Secure Firewall Management Center , and with Microsoft Active Directory.

Table 137: Passive Identity Agent port requirements

Port	Reason
443	Communicate with the Secure Firewall Management Center .
135	Communicate with Microsoft Active Directory using the MSRPC protocol.
9095	Communicate with each other using the UDP protocol.

History for the Passive Identity Agent

Table 138: History for the Passive Identity Agent

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Passive Identity Agent	7.6	7.1	This feature is introduced. Passive Identity Agent version 1.1 is compatible with 7.6.0 and later and adds the following: • You can use either FQDN, IPv4, or IPv6 to connect from the Passive Identity Agent to the Secure Firewall Management Center or Cisco Security Cloud Control. • Sends both IPv4 and IPv6 user sessions from Microsoft Active Directory (AD) to the Firewall Management Center. • You can zip troubleshooting logs. • When you start the Passive Identity Agent software, a list of prerequisites is displayed. The Passive Identity Agent identity source sends session data from Microsoft Active Directory (AD) to the Firewall Management Center. Passive identity agent software is supported on: • Microsoft AD server (Windows Server 2008 or later) • Microsoft AD domain controller (Windows Server 2008 or later) • Any client connected to the domain you want to monitor (Windows 8 or later)

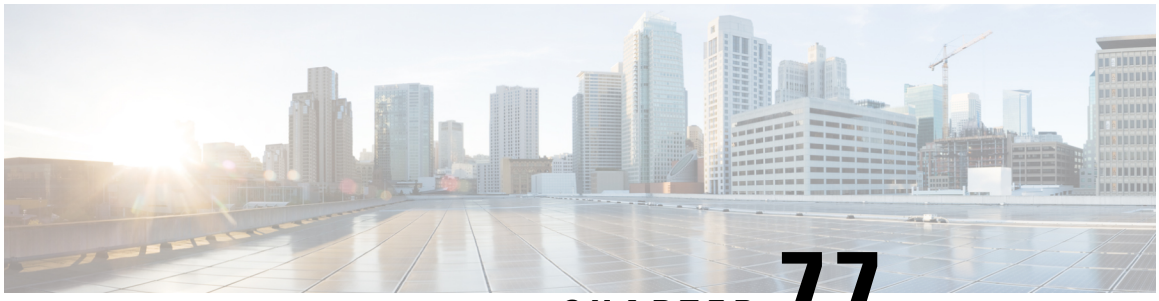

CHAPTER 77

User Control with ISE/ISE-PIC

The following topics discuss how to perform user awareness and user control with ISE/ISE-PIC:

- The ISE/ISE-PIC Identity Source, on page 2499
- License Requirements for ISE/ISE-PIC, on page 2501
- Requirements and Prerequisites for ISE/ISE-PIC, on page 2501
- ISE/ISE-PIC Guidelines and Limitations, on page 2501
- How to Configure ISE/ISE-PIC for User Control, on page 2504
- Configure ISE/ISE-PIC, on page 2506
- Ways to Configure the Cisco Identity Services Engine (Cisco ISE) Identity Source, on page 2512
- Troubleshoot the ISE/ISE-PIC or Cisco TrustSec Issues, on page 2521
- History for ISE/ISE-PIC, on page 2523

The ISE/ISE-PIC Identity Source

You can integrate your Cisco Identity Services Engine (ISE) or ISE Passive Identity Connector (ISE-PIC) deployment with the system to use ISE/ISE-PIC for passive authentication.

ISE/ISE-PIC is an authoritative identity source, and provides user awareness data for users who authenticate using Active Directory (AD), LDAP, RADIUS, or RSA. Additionally, you can perform user control on Active Directory users. ISE/ISE-PIC does not report failed login attempts or the activity of ISE Guest Services users.

In addition to user awareness and control, if you use ISE Cisco ISE to define and use security group tags (SGT) for classifying traffic in a Cisco TrustSec network, you can write access control rules that use SGT as both source and destination matching criteria. This enables you to block or allow access based on security group membership rather than IP addresses or network objects. For more information, see Configure Dynamic Attributes Conditions. Also see ISE/ISE-PIC Guidelines and Limitations, on page 2501.

Note The system does not parse IEEE 802.1x machine authentication but it *does* parse 802.1x user authentication. If you are using 802.1x with ISE, you must include user authentication. 802.1x machine authentication will not provide a user identity to the Firewall Management Center that can be used in policy.

For more information on Cisco ISE/ISE-PIC, see the Cisco Identity Services Engine Passive Identity Connector Administrator Guide or the Cisco Identity Services Engine Administrator Guide.

Note We strongly recommend you use the latest version of ISE/ISE-PIC to get the latest feature set and the most number of issue fixes.

Source and Destination Security Group Tag (SGT) Matching

If you use Cisco ISE to define and use security group tags (SGT) for classifying traffic in a Cisco TrustSec network, you can write access control rules that use SGT as both source and destination matching criteria. This enables you to block or allow access based on security group membership rather than IP addresses or network objects. For more information, see Configure Dynamic Attributes Conditions.

Matching on SGT tags provides the following benefits:

- The Firewall Management Center can subscribe to Security Group Tag eXchange Protocol (SXP) mappings from ISE.

  ISE uses SXP to propagate the IP-to-SGT mapping database to managed devices. When you configure Firewall Management Center to use an ISE server, you enable the option to listen to the SXP topic from ISE. This causes the Firewall Management Center to learn about the security group tags and mappings directly from ISE. The Firewall Management Center then publishes SGTs and mappings to managed devices.

  The SXP Topic receives security group tags based on static and dynamic mappings learned through the SXP protocol between ISE and other SXP compliant devices (like switches).

  You can create security group tags in ISE and assign host or network IP addresses to each tag. You can also assign SGTs to user accounts, and the SGT is assigned to the user's traffic. If the switches and routers in the network are configured to do so, these tags then get assigned to packets as they enter the network controlled by ISE, the Cisco TrustSec cloud.

  SXP is *not* supported by ISE-PIC.

- The Firewall Management Center and managed devices can learn about SGT mappings without deploying additional policy. (In other words, you can view connection events for SGT mappings without deploying an access control policy.)

- Supports Cisco TrustSec, which enables you to segment your network to protect critical business assets.

- When a managed device evaluates SGT as a traffic matching criteria for a rule, it uses the following priority:

  1. The source SGT tag defined in the packet, if any.

     For the SGT tag to be in the packet, the switches and routers in the network must be configured to add them. See the ISE documentation for information on how to implement this method.

     For the SGT tag to be in the packet, the switches and routers in the network must be configured to add them. See the ISE documentation for information on how to implement this method.

  2. The SGT assigned to the user session, as downloaded from the ISE session directory. The SGT can be matched to source or destination.

  3. The SGT-to-IP address mapping downloaded using SXP. If the IP address is in the range for an SGT, then the traffic matches the rule that uses the SGT. The SGT can be matched to source or destination.

Examples:

- In ISE, create an SGT tag named Guest Users and associate it with the 192.0.2.0/24 network.

  For example, you could use Guest Users as a source SGT condition in your access control rule and restrict access to certain URLs, web site categories, or networks from anyone who accesses your network.

- In ISE, create an SGT tag named Restricted Networks and associate it with the 198.51.100.0/8 network.

  For example, you could use Restricted Networks as a destination SGT rule condition and block access from Guest Users and other networks that have users who are not authorized to access the network.

Related Topics

ISE/ISE-PIC Guidelines and Limitations, on page 2501

License Requirements for ISE/ISE-PIC

Threat Defense License

Any

Requirements and Prerequisites for ISE/ISE-PIC

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

ISE/ISE-PIC Guidelines and Limitations

Use the guidelines discussed in this section when configuring ISE/ISE-PIC.

ISE/ISE-PIC Version and Configuration Compatibility

Your ISE/ISE-PIC version and configuration affects its integration and interaction with the Secure Firewall Management Center, as follows:

- We strongly recommend you use the latest version of ISE/ISE-PIC to get the latest feature set.
- Synchronize the time on the ISE/ISE-PIC server and the Secure Firewall Management Center. Otherwise, the system might perform user timeouts at unexpected intervals.

- To implement user control using ISE or ISE-PIC data, configure and enable a realm for the ISE server assuming the pxGrid persona as described in Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.
- Each Secure Firewall Management Center host name that connects to an ISE server must be unique; otherwise, the connection to one of the Secure Firewall Management Centers will be dropped.
- If you configure ISE/ISE-PIC to monitor a large number of user groups, the system might drop user mappings based on groups due to managed device memory limitations. As a result, rules with realm or user conditions might not perform as expected.

  For any device running version 6.7 or later, you can optionally use the **configure identity-subnet-filter** command to limit the subnets that the managed device monitors. For more information, see the Cisco Secure Firewall Threat Defense Command Reference.

  Alternatively, you can configure a network object and apply that object as an Identity Mapping Filter in the identity policy. See Create an Identity Policy, on page 2555.

For the specific versions of ISE/ISE-PIC that are compatible with this version of the system, see the Cisco Firepower Compatibility Guide.

IPv6 support

- Compatible versions of ISE/ISE-PIC version 2.x include support for IPv6-enabled endpoints.
- Version 3.0 (patch 2) and later of ISE/ISE-PIC enables IPv6 communication between ISE/ISE-PIC and the Firewall Management Center.

Approve clients in ISE

Before a connection between the ISE server and the Firewall Management Center succeeds, you must manually approve the clients in ISE. (Typically, there are two clients: one for the connection test and another for ISE agent.)

You can also enable **Automatically approve new accounts** in ISE as discussed in the chapter on Managing users and external identity sources in the *Cisco Identity Services Engine Administrator Guide*.

Unreachable sessions are removed

If a user session in ISE/ISE-PIC is reported as unreachable, the > prunes that session so another user with the same IP cannot match the unreachable user's identity rules.
You can control this behavior in ISE/ISE-PIC by going to **Providers** > **Endpoint Probes** and clicking one of the following:

- **Enabled** to cause ISE/ISE-PIC to monitor endpoint connections and therefore to cause the Secure Firewall Management Center prune a session from an unreachable user.
- **Disabled** to cause ISE/ISE-PIC to ignore endpoint connections.

Security Group Tags (SGT)

A Security Group Tag (SGT) specifies the privileges of a traffic source within a trusted network. Cisco ISE and Cisco TrustSec use a feature called Security Group Access (SGA) to apply SGT attributes to packets as they enter the network. These SGTs correspond to a user's assigned security group within ISE or TrustSec. If you configure ISE as an identity source, the Firepower System can use these SGTs to filter traffic.

Security Group Tags can be used both as source and destination matching criteria in access control rules.

Note To implement user control using only the ISE SGT attribute tag, you do not need to configure a realm for the ISE server. ISE SGT attribute conditions can be configured in policies with or without an associated identity policy.

Note In some rules, custom SGT conditions can match traffic tagged with SGT attributes that were *not* assigned by ISE. This is not considered user control, and works only if you are not using ISE/ISE-PIC as an identity source; see Custom SGT Conditions.

To match destination SGT tags in addition to source SGT tags, the following apply:

Required ISE version: 2.6 patch 6 or later, 2.7 patch 2 or later

Router support: Any Cisco router that supports SGT inline tagging over Ethernet. For more information, consult a reference such as the Cisco Group Based Policy Platform and Capability Matrix Release

Limitations:

- Quality of Service (QoS) policy uses source SGT matching only; it does *not* use destination SGT matching
- RA-VPN does not receive SGT mappings directly through RADIUS

ISE and High Availability

When the primary ISE/ISE-PIC server fails, the following occurs:

As a result of the integration with pxGrid v2, the Secure Firewall Management Center round-robins between both configured ISE hosts until one accepts the connection.

If the connection is lost, the Secure Firewall Management Center resumes round-robin attempts to the connected hosts.

Endpoint Location (or Location IP)

An Endpoint Location attribute is the IP address of the network device that used ISE to authenticate the user, as identified by ISE.

You must configure and deploy an identity policy to control traffic based on **Endpoint Location (Location IP)**.

ISE Attributes

Configuring an ISE connection populates the Secure Firewall Management Center database with ISE attribute data. You can use the following ISE attributes for user awareness and user control. This is not supported with ISE-PIC.

Endpoint Profile (or Device Type)

An Endpoint Profile attribute is the user's endpoint device type, as identified by ISE.

You must configure and deploy an identity policy to control traffic based on **Endpoint Profile (Device Type)**.

How to Configure ISE/ISE-PIC for User Control

You can use ISE/ISE-PIC in any of the following configurations:

- With a realm, identity policy, and associated access control policy.

  Use a realm to control *user* access to network resources in policy. You can still use ISE/ISE-PIC Security Group Tags (SGT) metadata in your policies.

- With an access control policy only. No realm or identity policy are necessary.

  Use this method to control network access using SGT metadata alone.

Related Topics

How to Configure ISE/ISE-PIC Without a Realm, on page 2504
How to Configure ISE/ISE-PIC for User Control Using a Realm, on page 2505

How to Configure ISE/ISE-PIC Without a Realm

This topic provides a high-level overview of tasks you must complete to configure ISE to be able to allow or block access to the network using SGT tags.

Procedure

	Command or Action	Purpose
Step 1	SGT matching: Enable SXP on ISE.	This enables the Secure Firewall Management Center to receive updates from ISE when SGT metadata changes.
Step 2	Create the ISE/ISE-PIC identity source.	The ISE/ISE-PIC identity source enables you to control user activity using Security Group Tags (SGT) provided by ISE/ISE-PIC. See Ways to Configure the Cisco Identity Services Engine (Cisco ISE) Identity Source, on page 2512.
Step 3	Create an access control rule.	The access control rule specifies an action to take (for example, allow or block) if traffic matches the rule criteria. You can use source and destination SGT metadata as matching criteria in the access control rule. See Introduction to Access Control Rules, on page 1773.
Step 4	Deploy the access control policy to managed devices.	Before your policy can take effect, it must be deployed to managed devices. See Deploy Configuration Changes, on page 259.

What to do next

Export Certificates from the ISE/ISE-PIC Server for Use in the Firewall Management Center, on page 2509

How to Configure ISE/ISE-PIC for User Control Using a Realm

Before you begin

This topic provides a high-level overview of tasks you must complete to configure ISE/ISE-PIC for user control and to be able to allow or block user or group access to the network. Users and groups can be stored in any server listed in Supported Servers for Realms, on page 2434.

Procedure

	Command or Action	Purpose
Step 1	Destination SGT only: Enable SXP on ISE.	This enables the Secure Firewall Management Center to receive updates from ISE when SGT metadata changes.
Step 2	Create a realm.	You must create a realm only to control access to the network by the users and groups you choose. See Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.
Step 3	Download users and groups, and enable the realm.	Downloading users and groups enables you to use them in access control rules. See Synchronize Users and Groups, on page 2445.
Step 4	Create the ISE/ISE-PIC identity source.	The ISE/ISE-PIC identity source enables you to control user activity using Security Group Tags (SGT) provided by ISE/ISE-PIC. See Ways to Configure the Cisco Identity Services Engine (Cisco ISE) Identity Source, on page 2512.
Step 5	Create an identity policy.	An identity policy is a container for one or more identity rules. See Create an Identity Policy, on page 2555.
Step 6	Create an identity rule.	An identity rule specifies how a realm is used to control access to the network by users and groups. See Create an Identity Rule, on page 2563.
Step 7	Associate the identity policy with an access control policy.	This enables the access control policy to use users and groups in the realm.
Step 8	Create an access control rule.	The access control rule specifies an action to take (for example, allow or block) if traffic

	Command or Action	Purpose
		matches the rule criteria. You can use source and destination SGT metadata as matching criteria in the access control rule. See Introduction to Access Control Rules, on page 1773.
Step 9	Deploy the access control policy to managed devices.	Before your policy can take effect, it must be deployed to managed devices. See Deploy Configuration Changes, on page 259.

What to do next

Export Certificates from the ISE/ISE-PIC Server for Use in the Firewall Management Center, on page 2509

Configure ISE/ISE-PIC

The following topics discuss how to configure the ISE/ISE-PIC server for use with identity policies in the Firewall Management Center.

The topics discuss how to:

- (Cisco ISE advanced configuration only.) Export certificates from the ISE/ISE-PIC server to authenticate with the Firewall Management Center.
- Publish SXP topics so the Firewall Management Center can be updated with Security Group Tags (SGT) on the ISE server.

Related Topics

Configure Security Groups and SXP Publishing in ISE, on page 2506
Export Certificates from the ISE/ISE-PIC Server for Use in the Firewall Management Center, on page 2509

Configure Security Groups and SXP Publishing in ISE

There is a lot of configuration that you must do in Cisco Identity Services Engine (ISE) to create the TrustSec policy and security group tags (SGT). Please look at the ISE documentation for more complete information on implementing TrustSec.

The following procedure picks out the highlights of the core settings you must configure in ISE for the Firewall Threat Defense device to be able to download and apply static SGT-to-IP address mappings, which can then be used for source and destination SGT matching in access control rules. See the ISE documentation for detailed information.

The screen shots in this procedure are based on ISE 2.4. The exact paths to these features might change in subsequent releases, but the concepts and requirements will be the same. Although ISE 2.4 or later is recommended, and preferably 2.6 or later, the configuration should work starting with ISE 2.2 patch 1.

Before you begin

You must have the ISE Plus license to publish SGT-to-IP address static mappings and to get user session-to-SGT mappings so that the Firewall Threat Defense device can receive them.

Procedure

Step 1 Choose **Work Centers** > **TrustSec** > **Settings** > **SXP Settings**, and select the **Publish SXP Bindings on PxGrid** option.

This option makes ISE send the SGT mappings out using SXP. You must select this option for the Firewall Threat Defense device to "hear" anything from listing to the SXP topic. This option must be selected for the Firewall Threat Defense device to get static SGT-to-IP address mapping information. It is not necessary if you simply want to use SGT tags defined in the packets, or SGTs that are assigned to a user session.

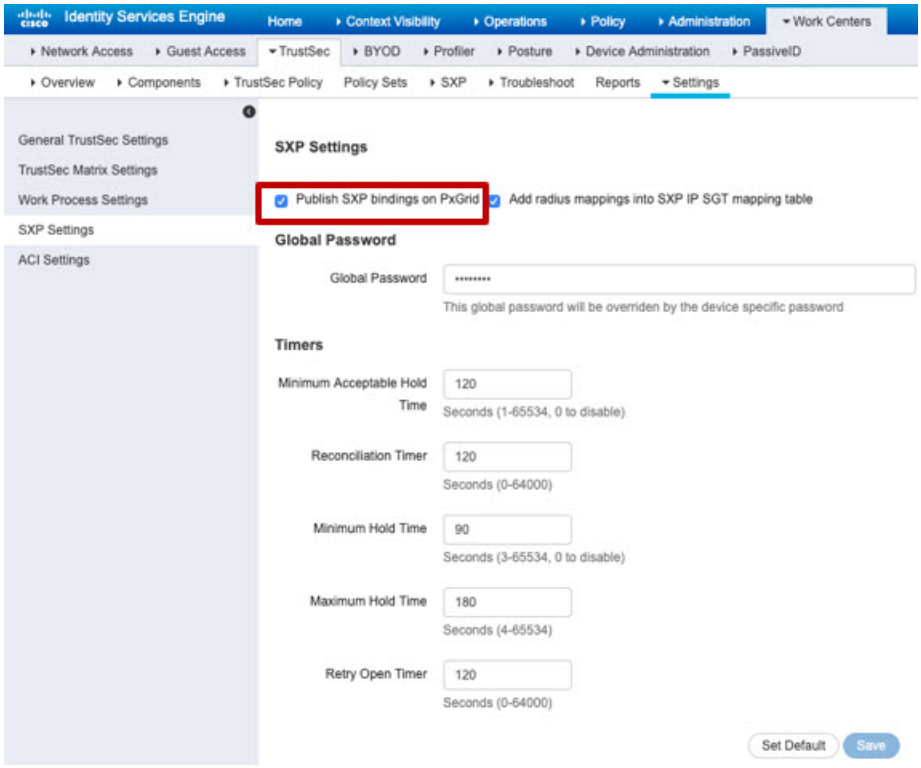

Step 2 Choose **Work Centers** > **TrustSec** > **SXP** > **SXP Devices**, and add a device.

This does not have to be a real device, you can even use the management IP address of the Firewall Threat Defense device. The table simply needs at least one device to induce ISE to publish the static SGT-to-IP address mappings. This step is not necessary if you simply want to use SGT tags defined in the packets, or SGTs that are assigned to a user session.

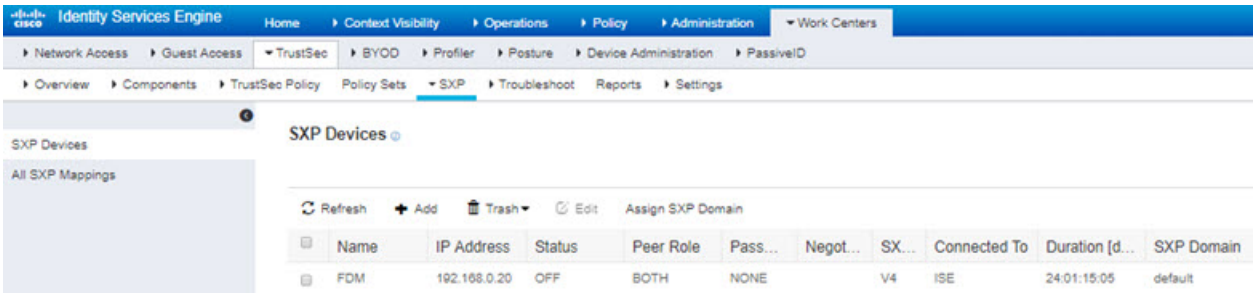

Step 3 Choose **Work Centers** > **TrustSec** > **Components** > **Security Groups** and verify there are security group tags defined. Create new ones as necessary.

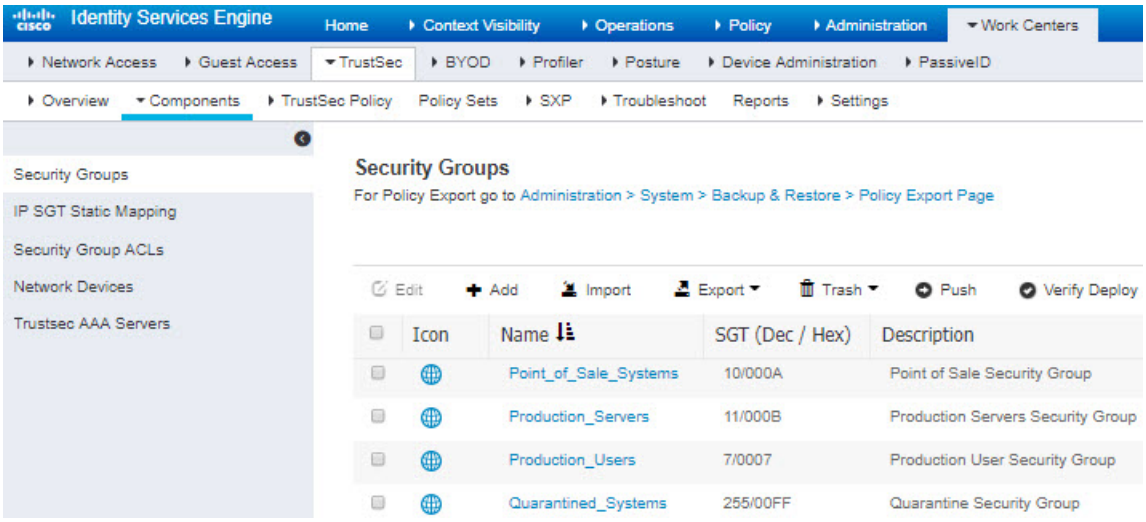

Step 4 Choose **Work Centers** > **TrustSec** > **Components** > **IP SGT Static Mapping** and map host and network IP addresses to the security group tags.

This step is not necessary if you simply want to use SGT tags defined in the packets, or SGTs that are assigned to a user session.

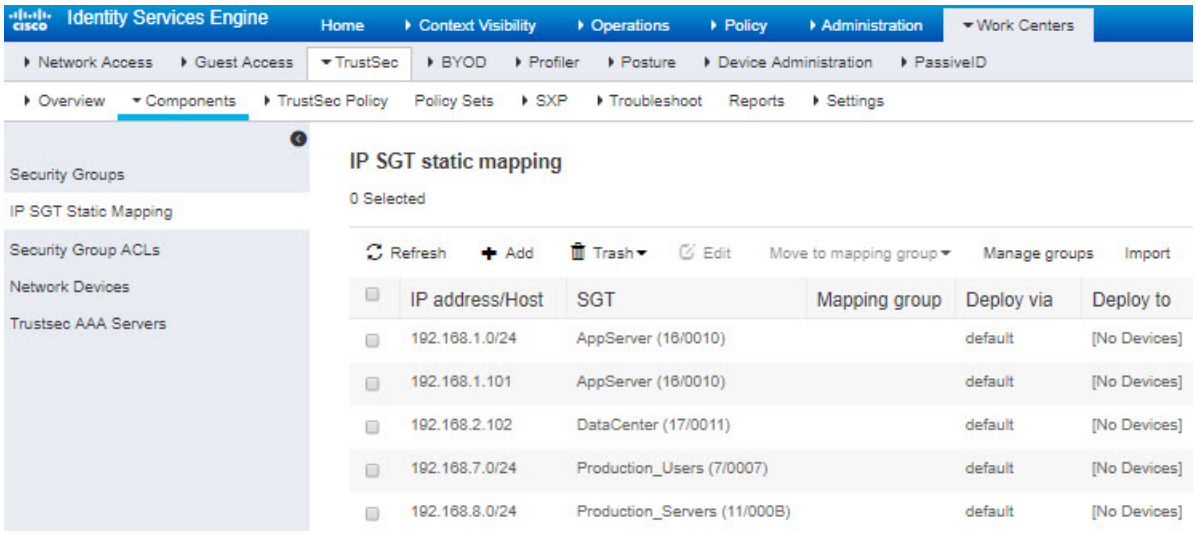

Export Certificates from the ISE/ISE-PIC Server for Use in the Firewall Management Center

The following sections discuss how to:

- Export system certificates from the ISE/ISE-PIC server.

  These certificates are required to securely connect to the ISE/ISE-PIC server. You might need to export one, or as many as three, certificates, depending on how your ISE system is set up:

  - One certificate for the pxGrid server
  - One certificate for the monitoring (MNT) server
  - One certificate, including the private key, for the pxGrid client (that is, the Firewall Management Center)

    Unlike the first two certificates, this is a self-signed certificate.

- Import these certificates into the Firewall Management Center:
  - pxGrid client certificate: internal certificate with key (**Objects** > **Object Management** > **PKI** > **Internal Certs**)
  - pxGrid server certificate: trusted CA (**Objects** > **Object Management** > **PKI** > **Trusted CAs**)
  - MNT certificate: trusted CA

Related Topics

Export a System Certificate
Import ISE/ISE-PIC Certificates, on page 2511

Export a System Certificate

You can export a system certificate or a certificate and its associated private key. If you export a certificate and its private key for backup purposes, you can reimport them later if needed.

Before you begin

To perform the following task, you must be a Super Admin or System Admin.

Procedure

Step 1 In the Cisco ISE GUI, click the **Menu** icon (☰) and choose **Administration** > **System** > **Certificates** > **System Certificates**.

Step 2 Check the check box next to the certificate that you want to export and click **Export**.

Step 3 Choose whether to export only the certificate, or the certificate and its associated private key.

Tip
We do not recommend exporting the private key that is associated with a certificate because its value may be exposed. If you must export a private key (for example, when you export a wildcard system certificate to be imported into the other Cisco ISE nodes for inter-node communication), specify an encryption password for the private key. You must specify this password while importing this certificate into another Cisco ISE node to decrypt the private key.

Step 4 Enter the password if you have chosen to export the private key. The password should be at least eight characters long.

Step 5 Click **Export** to save the certificate to the file system that is running your client browser.

If you export only the certificate, the certificate is stored in the PEM format. If you export both the certificate and private key, the certificate is exported as a .zip file that contains the certificate in the PEM format and the encrypted private key file.

Generate a Self-Signed Certificate

Add a new local certificate by generating a self-signed certificate. Cisco recommends that you only employ self-signed certificates for your internal testing and evaluation needs. If you plan to deploy Cisco ISE in a production environment, use CA-signed certificates whenever possible to ensure more uniform acceptance around a production network.

Note If you use a self-signed certificate and you want to change the hostname of your Cisco ISE node, log in to the administration portal of the Cisco ISE node, delete the self-signed certificate that has the old hostname, and generate a new self-signed certificate. Otherwise, Cisco ISE continues to use the self-signed certificate with the old hostname.

Before you begin

To perform the following task, you must be a Super Admin or System Admin.

Procedure

Step 1 In the Cisco ISE GUI, click the **Menu** icon (☰) and choose **Administration** > **System** > **Certificates** > **System Certificates**.

To generate a self-signed certificate from a secondary node, choose **Administration** > **System** > **Server Certificate**.

Step 2 In the ISE-PIC GUI, click the **Menu** icon (☰) and choose **Certificates** > **System Certificates**.

Step 3 Click **Generate Self Signed Certificate** and enter the details in the window displayed.

Step 4 Check the check boxes in the **Usage** area based on the service for which you want to use this certificate.

Step 5 Click **Submit** to generate the certificate.

To restart the secondary nodes, from the CLI, enter the following commands in the following order:

a) **application stop ise**
b) **application start ise**

Import ISE/ISE-PIC Certificates

This procedure is optional. You can also import the certificates automatically by providing login credentials to your ISE server as discussed in Ways to Configure the Cisco Identity Services Engine (Cisco ISE) Identity Source, on page 2512.

Import the certificates as followed:

- pxGrid client certificate: internal certificate with key (**Objects** > **Object Management** > **PKI** > **Internal Certs**)
- pxGrid server certificate: trusted CA (**Objects** > **Object Management** > **PKI** > **Trusted CAs**)
- MNT certificate: trusted CA

Before you begin

(Optional.) Export certificates from the ISE/ISE-PIC server as discussed in Export a System Certificate, on page 2510. The certificates and key must be present on the machine from which you log in to the Secure Firewall Management Center.

Procedure

Step 1 Log in to the Secure Firewall Management Center if you have not already done so.

Step 2 Click **Objects** > **Object Management** > **PKI** > **Internal Certs**.

Step 3 Click **Add Internal Cert**.

Step 4 Follow the prompts on your screen to import the certificate and private key.

Step 5 Click **Trusted CAs**.

Step 6 Click **Add Trusted CA**.

Step 7 Follow the prompts on your screen to import the pxGrid server certificate.

Step 8 Repeat the preceding steps, if necessary, to import the MNT server's trusted CA.

What to do next

Ways to Configure the Cisco Identity Services Engine (Cisco ISE) Identity Source, on page 2512

Ways to Configure the Cisco Identity Services Engine (Cisco ISE) Identity Source

You can configure the Cisco ISE identity source in any of the following ways:

- Quick configuration (new): *Cisco ISE only; not supported for ISE-PIC*. Enter the user name and password of a user in at least the External RESTful Services (ERS) Operator group. The Secure Firewall Management Center logs into the Cisco ISE Primary Authentication Node (PAN), downloads the certificates, and configures the identity source. You have the option of saving the ISE configuration for later use.

  For more information about groups, see the section on Cisco ISE Administrator Groups in the Cisco Identity Services Engine Administrator Guide.

  For details, see Quick Configuration, on page 2514.

- Advanced configuration (old): Same as previous Secure Firewall Management Center releases and works for both Cisco ISE and ISE-PIC. You must get certificates and other information from your ISE/ISE-PIC server as discussed in Configure ISE/ISE-PIC, on page 2506.

  For details about advanced configuration, see Cisco ISE Advanced Configuration, on page 2518.

Related Topics

About Cisco ISE Quick Configuration, on page 2512
Cisco ISE Advanced Configuration, on page 2518

About Cisco ISE Quick Configuration

Complete these tasks first: Prerequisites for ISE Quick Configuration, on page 2513.

You must have the following information:

- Fully qualified domain name or IP address of the policy administration node (PAN).
- User name and password of a user in at least the External RESTful Services (ERS) Operator group.

Note the following about SGT-to-IP mappings and SXP:

- To get all mappings that are defined in Cisco ISE, including SGT-to-IP address mappings published through SXP, use the procedure that follows. As an alternative, you have the following options:
  - To use the SGT information in the packets only, and not use mappings downloaded from Cisco ISE, skip the steps discussed in Create and Edit Access Control Rules, on page 1783. Note that in this case, you can use SGT tags as a source condition only; these tags will never match destination criteria.

- To use SGT in packets and user-to-IP-address/SGT mappings only, do not subscribe to the SXP topic in the Cisco ISE identity source, and do not configure ISE to publish SXP mappings. You can use this information for both source and destination matching conditions.

Related Topics

Prerequisites for ISE Quick Configuration, on page 2513
Quick Configuration, on page 2514
About Cisco ISE Quick Configuration, on page 2512
Cisco ISE Advanced Configuration, on page 2518

Prerequisites for ISE Quick Configuration

For ISE quick configuration to succeed, you mst do both of the following as an ISE administrator:

- Enable External RESTful Services (ERS) (read/write), Open API (read/write), and *disable* the Cross-Site Request Forgery (CSRF) check

  Go to > **Administration** > **System** > **Settings** > **API Settings** > **API Service Settings**

  For more information, see Enable API Service in the Cisco Identity Services Engine Administrator Guide.

  The following figure shows an example.

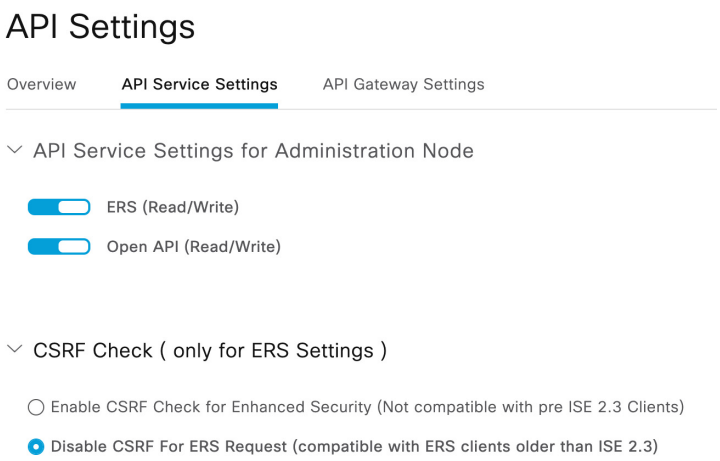

- Slide **ERS (Read/Write)** to **Enabled**.
- Enable the certificate authority

  Go to > **Administration** > **System** > **Certificates** > **Certificate Authority** > **Internal CA Settings**

  For more information, see Internal CA Settings in the Cisco Identity Services Engine Administrator Guide.

  The following figure shows an example.

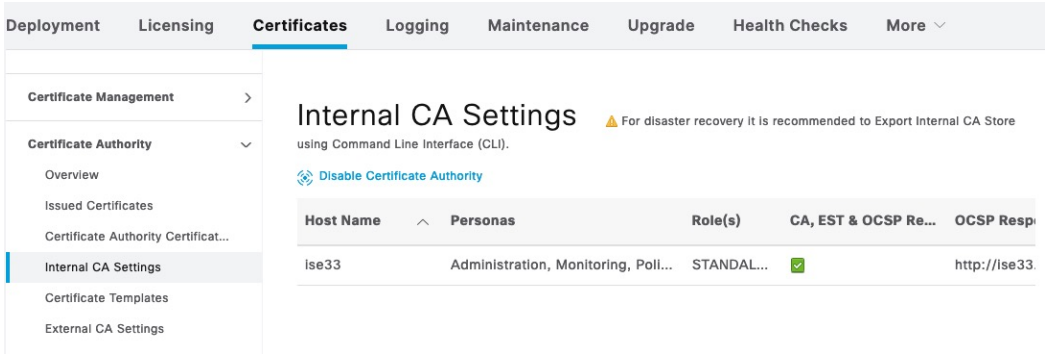

Related Topics

About Cisco ISE Quick Configuration, on page 2512
Cisco ISE Advanced Configuration, on page 2518

Quick Configuration

This task discusses how to configure Cisco ISE (but not ISE-PIC) by entering a user name and password. The Secure Firewall Management Center then logs in to ISE and downloads the necessary certificates to authenticate the two applications.

Firewall Threat Defense Feature History:

7.6—This feature is introduced.

Before you begin

See the following topics first:

- About Cisco ISE Quick Configuration, on page 2512
- Prerequisites for ISE Quick Configuration, on page 2513

Procedure

Step 1 Log in to the Secure Firewall Management Center.

Step 2 Click **Integration** > **Other Integrations** > **Identity Sources**.

Step 3 Click **Identity Services Engine** for the **Service Type** to enable the ISE connection.

Note
To disable the connection, click **None**.

Step 4 Click **Quick Configuration (New)**.

Step 5 In the **Primary PAN FQDN/IP Address** field, enter the fully qualified domain name or IP address of the policy administration node (PAN). Do *not* enter a scheme (such as **https://**).

Step 6 In the **Username** field, enter the user name of a user in at least the ERS Operator group.

For more information about groups, see the section on Cisco ISE Administrator Groups in the Cisco Identity Services Engine Administrator Guide.

Step 7 In the **Password** field, enter the user's password.

Step 8 (Optional.) Enter an **ISE Network Filter** using CIDR block notation.

Step 9 In the Subscribe To section, check the following:

- **Session Directory Topic** to receive ISE user session information from the ISE server.
- **SXP Topic** to receive updates to SGT-to-IP mappings when available from the ISE server. This option is required to use destination SGT tagging in access control rules.

Step 10 To test the connection, click **Test**.

Connection and integration results are discussed in Cisco Identity Services Engine (Cisco ISE) Quick Configuration Results, on page 2515.

Step 11 (Optional.) After a successful trest, click **Save this Config** at the top of the page to save the configuration on the Secure Firewall Management Center.

What to do next

- Specify users to control and other options using an identity policy as described in Create an Identity Policy, on page 2555.
- Associate the identity rule with an access control policy, which filters and optionally inspects traffic, as discussed in Associating Other Policies with Access Control, on page 1762.
- Use Security Group Tags (SGT) from Cisco ISE as dynamic attributes in access control policies.

  For more information, see Configure Dynamic Attributes Conditions.
- Deploy your identity and access control policies to managed devices as discussed in Deploy Configuration Changes, on page 259.
- Monitor user activity as discussed in *Using Workflows* in the Cisco Secure Firewall Management Center Administration Guide.

Related Topics

About Cisco ISE Quick Configuration, on page 2512
Cisco ISE Advanced Configuration, on page 2518

Cisco Identity Services Engine (Cisco ISE) Quick Configuration Results

Initial information page for an untrusted certificate

When you enter required information in the **Quick Configuration (New)** tab page, the system first checks to see if the certificate was signed by a trusted authority. If not, a page similar to the following is displayed.

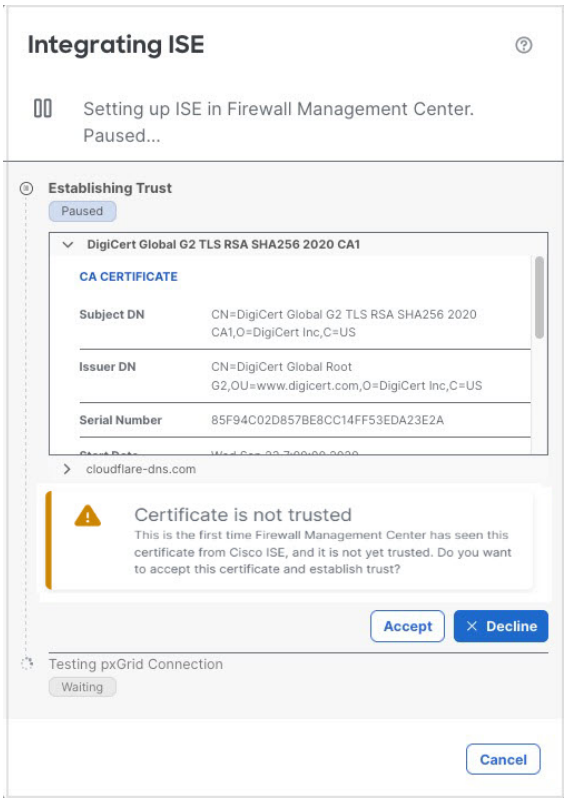

You have the following options:

- Click **Accept** for the > to trust the certificate and continue authenticating with ISE.
- Click **Decline** to quit without authenticating with ISE.

  You can then verify the ISE configuration on the **Quick Configuration (New)** tab page and try again.

Successful ISE integration

If the Secure Firewall Management Center was able to authenticate successfully with the information you provided, a page similar to the following is displayed.

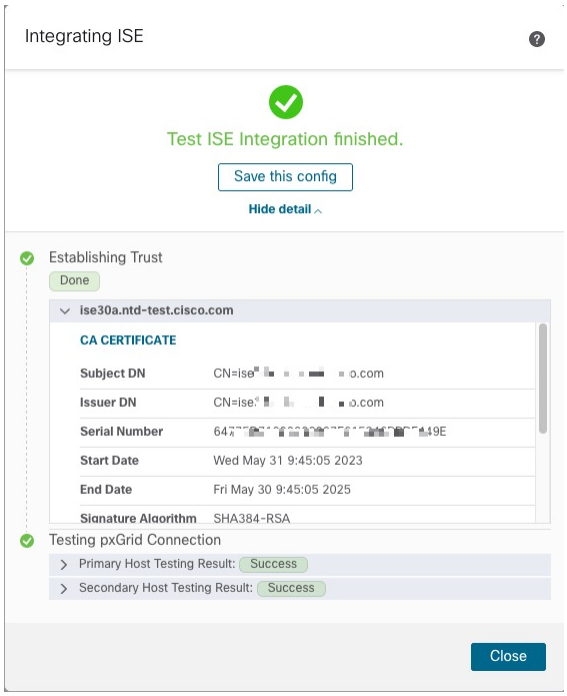

Click **Save this Config** to save it and be able to use the identity source in the future.

Optionally expand one of the lists under Testing pxGrid Connection to display certificates downloaded from those sources.

Unsuccessful ISE integration

If the Secure Firewall Management Center fails to authenticate with ISE for any reason, a page similar to the following is displayed.

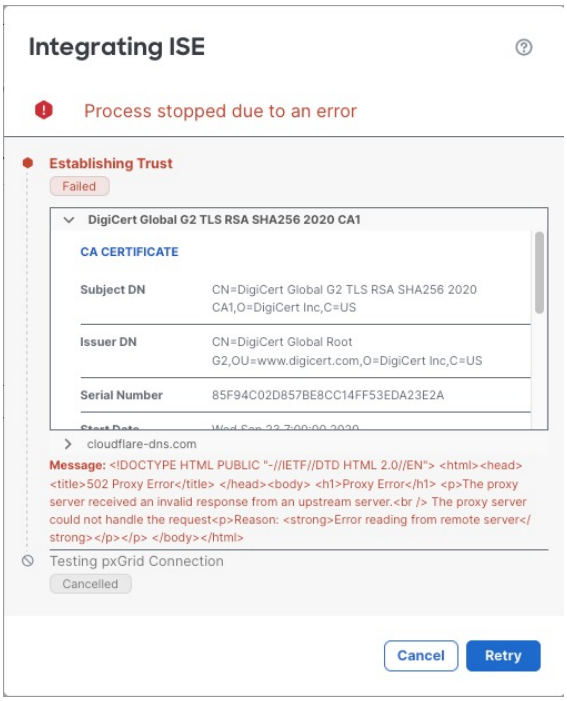

The error messages are displayed in red text.

You have the following options:

- If the error is temporary (for example, a temporarily network issue), click **Retry**.
- To try again with different ISE credentials, click **Cancel**.

Related Topics

About Cisco ISE Quick Configuration, on page 2512
Cisco ISE Advanced Configuration, on page 2518

Cisco ISE Advanced Configuration

The following procedure discusses how to configure the ISE/ISE-PIC identity source. You must be in the global domain to perform this task.

Before you begin

- To get user sessions from a Microsoft Active Directory Server or supported LDAP server, configure and enable a realm for the Cisco ISE server, assuming the pxGrid persona, as discussed in Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.
- To get all mappings that are defined in Cisco ISE, including SGT-to-IP address mappings published through SXP, use the procedure that follows. As an alternative, you have the following options:
  - To use the SGT information in the packets only, and not use mappings downloaded from Cisco ISE, skip the steps discussed in Create and Edit Access Control Rules, on page 1783. Note that in this case, you can use SGT tags as a source condition only; these tags will never match destination criteria.

- To use SGT in packets and user-to-IP-address/SGT mappings only, do not subscribe to the SXP topic in the Cisco ISE identity source, and do not configure ISE to publish SXP mappings. You can use this information for both source and destination matching conditions.

- (Advanced configuration only.) Export certificates from the ISE/ISE-PIC server and optionally import them into the Secure Firewall Management Center as discussed in Export Certificates from the ISE/ISE-PIC Server for Use in the Firewall Management Center, on page 2509.
- To publish SXP topics so the Secure Firewall Management Center can be updated with Security Group Tags (SGT) on the ISE server, see Configure ISE/ISE-PIC, on page 2506.

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Integration** > **Other Integrations** > **Identity Sources**.

Step 3 Click **Identity Services Engine** for the **Service Type** to enable the ISE connection.

Note
To disable the connection, click **None**.

Step 4 Enter a **Primary Host Name/IP Address** and, optionally, a **Secondary Host Name/IP Address**.

Step 5 Click the appropriate certificate authorities from the **pxGrid Server CA** and **MNT Server CA** lists, and the appropriate certificate from the **pxGrid Client Certificate** list. You can also click **Add** (+) to add a certificate.

Note
The **pxGrid Client Certificate** must include the **clientAuth** extended key usage value, or it must not include any extended key usage values.

Step 6 (Optional.) Enter an **ISE Network Filter** using CIDR block notation.

Step 7 In the Subscribe To section, check the following:

- **Session Directory Topic** to receive ISE user session information from the ISE server.
- **SXP Topic** to receive updates to SGT-to-IP mappings when available from the ISE server. This option is required to use destination SGT tagging in access control rules.

Step 8 To test the connection, click **Test**.

If the test fails, click **Additional Logs** for more information about the connection failure.

Note
If you set up load balancing with ISE, the node that responds to the test should be the currently active node. For more information, see Cisco pxGrid Node.

What to do next

- Specify users to control and other options using an identity policy as described in Create an Identity Policy, on page 2555.

- Associate the identity rule with an access control policy, which filters and optionally inspects traffic, as discussed in Associating Other Policies with Access Control, on page 1762.
- Use Security Group Tags (SGT) from Cisco ISE as dynamic attributes in access control policies.

  For more information, see Configure Dynamic Attributes Conditions.
- Deploy your identity and access control policies to managed devices as discussed in Deploy Configuration Changes, on page 259.
- Monitor user activity as discussed in *Using Workflows* in the Cisco Secure Firewall Management Center Administration Guide.

Related Topics

About Cisco ISE Quick Configuration, on page 2512
Cisco ISE Advanced Configuration, on page 2518

ISE/ISE-PIC Configuration Fields

The following fields are used to configure a connection to /ISE-PIC.

Primary and Secondary Host Name/IP Address

The hostname or IP address for the primary and, optionally, the secondary pxGrid ISE servers.

The ports used by the host names you specify must be reachable by both ISE and the Secure Firewall Management Center.

pxGrid Server CA

The trusted certificate authority for the pxGrid framework. If your deployment includes a primary and a secondary pxGrid node, the certificates for both nodes must be signed by the same certificate authority.

MNT Server CA

The trusted certificate authority for the ISE certificate when performing bulk downloads. If your deployment includes a primary and a secondary MNT node, the certificates for both nodes must be signed by the same certificate authority.

pxGrid Client Certificate

The internal certificate and key that the Secure Firewall Management Center must provide to /ISE-PIC to connect to /ISE-PIC or to perform bulk downloads.

Note The **pxGrid Client Certificate** must include the clientAuth extended key usage value, or it must not include any extended key usage values.

ISE Network Filter

An optional filter you can set to restrict the data that ISE reports to the Secure Firewall Management Center. If you provide a network filter, ISE reports data from the networks within that filter. You can specify a filter in the following ways:

- Leave the field blank to specify **`any`**.
- Enter a single IPv4 address block using CIDR notation.

- Enter a list of IPv4 address blocks using CIDR notation, separated by commas.

Note This version of the system does not support filtering using IPv6 addresses, regardless of your ISE version.

Subscribe to:

Session Directory Topic: Check this box to subscribe to user session information from the ISE server. Includes SGT and endpoint metadata.

SXP Topic: Check this box to subscribe to SXP mappings from the ISE server.

Related Topics

Trusted Certificate Authority Objects, on page 1409
Internal Certificate Objects, on page 1412
About Cisco ISE Quick Configuration, on page 2512
Cisco ISE Advanced Configuration, on page 2518

Troubleshoot the ISE/ISE-PIC or Cisco TrustSec Issues

Troubleshoot Cisco TrustSec issues

A device interface can be configured to propagate Security Group Tags (SGTs) either from ISE/ISE-PIC or from a Cisco device on the network (referred to as Cisco TrustSec.) On the device management page (**Devices** > **Device Management**), the **Propagate Security Group Tag** check box for an interface is disabled by default. If you want the interface to propagate TrustSec data, check the box.

Troubleshoot ISE/ISE-PIC issues

For other related troubleshooting information, see Troubleshoot Realms and User Downloads, on page 2456 and Troubleshoot User Control, on page 2573.

If you experience issues with the ISE or ISE-PIC connection, check the following:

- The pxGrid Identity Mapping feature in ISE must be enabled before you can successfully integrate ISE with the system.
- When the primary server fails, you must manually promote the secondary to primary; there is no automatic failover.
- Before a connection between the ISE server and the Firewall Management Center succeeds, you must manually approve the clients in ISE. (Typically, there are two clients: one for the connection test and another for ISE agent.)

  You can also enable **Automatically approve new accounts** in ISE as discussed in the chapter on Managing users and external identity sources in the Cisco Identity Services Engine Administrator Guide.
- The **pxGrid Client Certificate** must include the **clientAuth** extended key usage value, or it must not include any extended key usage values.
- The time on your ISE server must be synchronized with the time on the Secure Firewall Management Center. If the appliances are not synchronized, the system may perform user timeouts at unexpected intervals.

- If your deployment includes a primary and a secondary pxGrid node,
  - The certificates for both nodes must be signed by the same certificate authority.
  - The ports used by the host name must be reachable by both the ISE server and by the Secure Firewall Management Center.
- If your deployment includes a primary and a secondary MNT node, the certificates for both nodes must be signed by the same certificate authority.

To exclude subnets from receiving user-to-IP and Security Group Tag (SGT)-to-IP mappings from ISE, use the **configure identity-subnet-filter** {**add** | **remove**} command. You should typically do this for lower-memory managed devices to prevent Snort identity health monitor memory errors.

If you experience issues with user data reported by ISE or ISE-PIC, note the following:

- After the system detects activity from an ISE user whose data is not yet in the database, the system retrieves information about them from the server. Activity seen by the ISE user is *not* handled by access control rules, and is *not* displayed in the web interface until the system successfully retrieves information about them in a user download.
- You cannot perform user control on ISE users who were authenticated by an LDAP, RADIUS, or RSA domain controller.
- The Secure Firewall Management Center does not receive user data for ISE Guest Services users.
- If ISE monitors the same users as TS Agent, the Secure Firewall Management Center prioritizes the TS Agent data. If the TS Agent and ISE report identical activity from the same IP address, only the TS Agent data is logged to the Secure Firewall Management Center.
- Your ISE version and configuration impact how you can use ISE in the system. For more information, see The ISE/ISE-PIC Identity Source, on page 2499.
- If you have Secure Firewall Management Center high availability configured and the primary fails, see the section on ISE and high availability in ISE/ISE-PIC Guidelines and Limitations, on page 2501.
- ISE-PIC does not provide ISE attribute data.
- ISE-PIC cannot perform ISE ANC remediations.
- Active FTP sessions are displayed as the **Unknown** user in events. This is normal because, in active FTP, the server (not the client) initiates the connection and the FTP server should not have an associated user name. For more information about active FTP, see RFC 959.

If you experience issues with supported functionality, see The ISE/ISE-PIC Identity Source, on page 2499 for more information about version compatibility.

ISE/ISE-PIC user timeout

If you're setting up ISE/ISE-PIC without a realm, be aware there is a user session timeout that affects how users are seen by the Secure Firewall Management Center. For more information, see Realm Fields, on page 2436.

History for ISE/ISE-PIC

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Quick configuration	7.6	7.6.0	You can optionally configure ISE using only the user name and password of a user in the External RESTful Services (ERS) Operator group. (*This feature is available for ISE only, not ISE-PIC.*) **Upgrade impact**. Any ISE or ISE-PIC identity source you created before upgrade is still available on the **Advanced Configuration (Old)** tab page. Quick configuration affects only new ISE identity sources created after the upgrade. New/modified screens: **Integration** > **Other Integrations** > **Identity Sources** > **Identity Services Engine**. There are two tab pages: **Quick Configuration (New)** and **Advanced Configuration (Old)**. New/modified CLI commands: none
pxGrid 2.0 is the default for supported ISE/ISE-PIC versions	6.7.0	6.7.0	Note the following: • Supported ISE/ISE-PIC versions: 2.6 patch 6 or later, 2.7 patch 2 or later • Adaptive Network Control (ANC) policies replace Endpoint Protection Service (EPS) remediations. If you have EPS policies configured in the Firewall Management Center, you must migrate them to use ANC.
Optionally exclude subnets from receiving user-to-IP and Security Group Tag (SGT)-to-IP mappings from ISE. You should typically do this for lower-memory managed devices to prevent Snort identity health monitor memory errors.	6.7.0	6.7.0	New command: **configure identity-subnet-filter** {**add** | **remove**}

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Destination Security Group Tag matching (SGT)	6.5.0	6.5.0	Feature introduced. Enables you to use ISE SGT tags for both source and destination matching criteria in access control rules. SGT tags are tag-to-host/network mappings obtained by ISE. New/modified screens: • New options to configure Destination SGT matching: **System** > **Integration** > **Identity Sources** > **ISE/ISE-PIC** • **Session Directory Topic**: Subscribe to ISE user session information. • **SXP Topic**: Subscribe to SGT tag updates on the ISE server. • New and renamed columns in **Analysis** > **Connections** > **Events** • Renamed: Security Groups Tags renamed to Source SGT • New: Destination SGT
Integration with ISE-PIC	6.2.1	6.2.1	You can now use data from ISE-PIC.
SGT tags for user control.	6.2.1	6.2.0	You no longer need to create a realm or identity policy to perform user control based on ISE Security Group Tag (SGT) data.
Integration with ISE.	6.0	6.0	Feature introduced. By subscribing to Cisco's Platform Exchange Grid (PxGrid), the Firepower Management Center can download additional user data, device type data, device location data, and Security Group Tags (SGTs) —a method used by ISE to provide network access control).

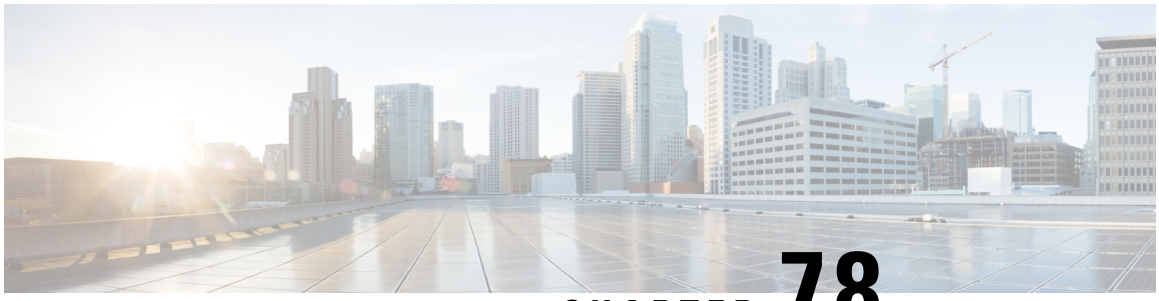

CHAPTER 78

User Control with Captive Portal

- The Captive Portal Identity Source, on page 2525
- License Requirements for Captive Portal, on page 2526
- Requirements and Prerequisites for Captive Portal, on page 2526
- Captive Portal Guidelines and Limitations, on page 2526
- How to Configure the Captive Portal for User Control, on page 2529
- Troubleshoot the Captive Portal Identity Source, on page 2541
- History for Captive Portal, on page 2543

The Captive Portal Identity Source

Captive portal is one of the authoritative identity sources supported by the system. Captive portal is an active authentication method where users authenticate onto the network using a managed device. (RA-VPN is another type of active authentication.) Active authentication differs from passive authentication in that the user is presented with a login page by the managed device, whereas passive authentication queries the authentication realm (for example, Microsoft AD) to authenticate the user.

You typically use captive portal to require authentication to access the internet or to access restricted internal resources; you can optionally configure guest access to resources. After the system authenticates captive portal users, it handles their user traffic according to access control rules. Captive portal performs authentication on HTTP and HTTPS traffic only.

Note To use a Microsoft Azure AD (SAML) realm for captive portal, see Create a Microsoft Azure AD (SAML) Realm for Active Authentication (Captive Portal), on page 2412.

Note HTTPS traffic must be decrypted before captive portal can perform authentication.

Captive portal also records failed authentication attempts. A failed attempt does not add a new user to the list of users in the database. The user activity type for failed authentication activity reported by captive portal is **Failed Auth User**.

The authentication data gained from captive portal can be used for user awareness and user control.

Related Topics

How to Configure the Captive Portal for User Control, on page 2529

About Hostname Redirect

(Snort 3 only.) An active authentication identity rule redirects to the captive portal port using its configured interface. Because the redirect is typically done to an IP address, the user gets an untrusted certificate error and because this behavior is similar to a man-in-the-middle attack, users might be reluctant to accept the untrusted certificate.

To avoid this problem, you can configure the captive portal to use the managed device's fully-qualified domain name (FQDN). With a properly configured certificate, users will not get an untrusted certificate error, and the authentication will be more seamless and appear to be more secure.

Related Topics

Redirect to Host Name Network Rule Conditions, on page 2558

License Requirements for Captive Portal

Threat Defense License

Any

Requirements and Prerequisites for Captive Portal

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Captive Portal Guidelines and Limitations

When you configure and deploy captive portal in an identity policy, users from specified realms authenticate using Firewall Threat Defense to access your network.

Note When a remote access VPN user has already actively authenticated through a managed device acting as a secure gateway, captive portal active authentication will not occur, even if configured in an identity policy.

Captive portal and policies

You configure captive portal in your identity policy and invoke active authentication in your identity rules. Identity policies are associated with access control policies and access control policies define access to resources in the network. For example, you might exclude users in the US-West/Finance group to access Engineering servers or you can prohibit users from accessing nonsecure applications on the network.

You configure some captive portal identity policy settings on the identity policy's **Active Authentication** tab page and configure the rest in the identity rule associated with the access control policy.

An active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive or VPN identity cannot be established** selected. In each case the system transparently enables or disables TLS/SSL decryption, which restarts the Snort process.

Caution Adding the first or removing the last active authentication rule when TLS/SSL decryption is disabled (that is, when the access control policy does not include a decryption policy) restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

When the captive portal authenticates users that match an identity rule, any user in a Microsoft Active Directory or LDAP group that has not been downloaded is identified as Unknown. To avoid users being identified as Unknown, configure the realm or realm sequence to download users in all groups you expect to authenticate with captive portal. Unknown users are handled according to the associated access control policy; if the access control policy is configured to block Unknown users, these users are blocked.

To make sure the system downloads all users in a realm or realm sequence, make sure the groups are in the Available Groups list in the realm's configuration.

For more information about synchronizing users and groups, see Synchronize Users and Groups, on page 2445.

Routed interface required

Captive portal active authentication can be performed only by a device with a routed interface configured. If you are configuring an identity rule for captive portal and your captive portal device contains inline and routed interfaces, you must configure interface rule conditions in the access control policy to target only the routed interfaces on the device.

If the identity policy associated with your access control policy contains one or more captive portal identity rules and you deploy the policy on the Secure Firewall Management Center that manages one or more devices with routed interfaces configured, the policy deployment succeeds and the routed interfaces perform active authentication.

Required certificate and certificate authorities

Before you can use the captive portal for user control and awareness, you must have all of the following:

- To authenticate with Microsoft AD, export the server's root certificate and import it into the Secure Firewall Management Center as a trusted CA certificate.
- An internal certificate object for authenticating with the managed device to which the identity policy is deployed.
- An internal certificate authority for the required decryption rule.

You can create the internal certificate and internal certificate authority at the time you create a decryption policy.

Captive portal requirements and limitations

Note the following requirements and limitations:

- Captive portal does not support HTTP/3 QUIC connections.
- The system supports up to 20 captive portal logins per second.
- There is a maximum five minute limit between failed login attempts for a failed login attempt to be counted toward the count of maximum login attempts. The five minute limit is not configurable.

  (Maximum login attempts are displayed in connection events: **Analysis** > **Connections**.)

  If more than five minutes elapse between failed logins, the user is redirected to captive portal for authentication and will not be designated a failed login user or a guest user, and will not be reported to the Secure Firewall Management Center.
- Captive portal does not negotiate TLS v1.0 connections.

  Only TLS v1.1, v1.2, and TLS 1.3 connections are supported.
- To securely log a user out of captive portal and prevent the user from logging in again, an administrator can close their session: **Analysis** > **Users heading** > **Active Sessions**. If the user closes their browser, the browser does not reauthenticate automatically.
- If a realm is created for a parent domain and the managed device detects a login to a child of that parent domain, the user's subsequent logout is not detected by the managed device.
- Your access control rule must allow traffic destined for the IP address and port of the device you plan to use for captive portal.
- To perform captive portal active authentication on HTTPS traffic, you must use a decryption policy to decrypt the traffic from the users you want to authenticate. You cannot decrypt the traffic in the connection between a captive portal user's web browser and the captive portal daemon on the managed device; this connection is used to authenticate the captive portal user.
- To limit the amount of non-HTTP or HTTPS traffic that is allowed through the managed device, you should enter typical HTTP and HTTPS ports in the identity policy's **Ports** tab page.

  The managed device changes a previously unseen user from **`Pending`** to **`Unknown`** when it determines that the incoming request does not use the HTTP or HTTPS protocol. As soon as the managed device changes a user from **`Pending`** to another state, access control, Quality of Service, and decryption policies can be applied to that traffic. If your other policies don't permit non-HTTP or HTTPS traffic, configuring ports on the captive portal identity policy can prevent undesired traffic from being allowed through the managed device.

Kerberos prerequisites

If you're using Kerberos authentication, the managed device's host name must be less than 15 characters (it's a NetBIOS limitation set by Windows); otherwise, captive portal authentication fails. You set the managed device host name when you set up the device. For more information, see an article like this one on the Microsoft documentation site: Naming conventions in Active Directory for computers, domains, sites, and OUs.

DNS must return a response of 64KB or less to the hostname; otherwise, the AD connection test fails. This limit applies in both directions and is discussed in RFC 6891 section-6.2.5.

How to Configure the Captive Portal for User Control

Before you begin

To use the captive portal for active authentication, you must set up an LDAP realm; or a Microsoft AD realm or realm sequence; Microsoft Azure AD (SAML) realm; access control policy; an identity policy; a decryption policy; and associate the identity and decryption policies with the same access control policy. Finally, you must deploy the policies to managed devices. This topic provides a high-level summary of those tasks.

Note To use a Microsoft Azure AD (SAML) realm as a captive portal, see How to Create a Microsoft Azure AD (SAML) Realm for Active Authentication (Captive Portal), on page 2403.

Perform the following tasks first:

- Confirm that your Secure Firewall Management Center manages one or more devices with a *routed* interface configured.
- To use encrypted authentication with the captive portal, either create a PKI object for the authenticating managed device or have your certificate data and key available on the machine from which you're accessing the Secure Firewall Management Center. To create a PKI object, see PKI, on page 1404.

Procedure

Step 1 Create and enable an LDAP realm; or a Microsoft AD realm and optionally realm sequence as discussed in the following topics:

- Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426
- Synchronize Users and Groups, on page 2445

To make sure the system downloads all users in a realm or realm sequence, make sure the groups are in the Available Groups list in the realm's configuration.

For more information, see Synchronize Users and Groups, on page 2445.

Step 2 Get required certificates and certificate authorities.
You must have all of the following:

- To authenticate with Microsoft AD, export the server's root certificate and import it into the Secure Firewall Management Center as a trusted CA certificate.
- An internal certificate object for authenticating with the managed device to which the identity policy is deployed.
- An internal certificate authority for the required decryption rule.

You can create the internal certificate and internal certificate authority at the time you create a decryption policy.

Step 3 Create a network object with an associated trusted certificate authority.

See Configure the Captive Portal Part 1: Create a Network Object, on page 2531.

Step 4 Create identity policy with an active authentication rule.

The identity policy enables selected users in your realm access resources after authenticating with the captive portal.

For more information, see Configure the Captive Portal Part 2: Create an Identity Policy and Active Authentication Rule, on page 2532.

Step 5 Configure an access control rule for the captive portal that allows traffic on the captive portal port (by default, TCP 885).

You can choose any available TCP port for the captive portal to use. Whatever your choice, you must create a rule that allows traffic on that port.

For more information, see Configure the Captive Portal Part 3: Create a TCP Port Access Control Rule, on page 2534.

Step 6 Add another access control rule to allow users in the selected realm or realm sequence to access resources using the captive portal.

For more information, see Configure the Captive Portal Part 4: Create a User Access Control Rule, on page 2535.

Step 7 Configure a decryption policy with a **Decrypt - Resign** rule for the **Unknown** user so captive portal users can access web pages using the HTTPS protocol.

The captive portal can authenticate users only if the HTTPS traffic is decrypted before the traffic is sent to the captive portal. The captive portal itself is seen by the system as the **Unknown** user.

Captive Portal Example: Create a Decryption Policy with an Outbound Rule, on page 2536

Step 8 Associate the identity and decryption policies with the access control policy from step 3.

This final step enables the system to authenticate users with the captive portal.

For more information, see Configure Captive Portal Part 6: Associate Identity and Decryption Policies with the Access Control Policy, on page 2538.

What to do next

See Configure the Captive Portal Part 1: Create a Network Object, on page 2531.

Related Topics

Exclude Applications from Captive Portal, on page 2540
PKI, on page 1404
Troubleshoot the Captive Portal Identity Source, on page 2541
Snort Restart Scenarios, on page 252

Configure the Captive Portal Part 1: Create a Network Object

This task discusses how to start configuring the captive portal as an identity source.

Before you begin

(Snort 3 only.) Create a fully-qualified host name (FQDN) using your DNS server and upload the Firewall Threat Defense's internal certificate to the Firewall Management Center. You can consult a resource such as this one if you've never done it before. Specify the IP address of a routed interface on one of the devices managed by your Firewall Management Center.

For more information about the network object, see Redirect to Host Name Network Rule Conditions, on page 2558.

Procedure

Step 1 If you haven't already done so, log in to your Firewall Management Center.

Step 2 Click **Objects** > **Object Management**.

Step 3 Expand **PKI**.

Step 4 Click **Internal Certs**.

Step 5 Click **Add Internal Cert**.

Step 6 In the **Name** field, enter a name to identify the internal cert (for example, **`MyCaptivePortal`**).

Step 7 In the **Certificate Data** field, either paste the certificate or use the **Browse** button to locate it.

The certificate Common Name must exactly match the FDQN with which you want captive portal users to authenticate.

Step 8 In the **Key** field, either paste the certificate's private key or use the **Browse** button to locate it.

Step 9 If the certificate is encrypted, select the **Encrypted** check box and enter the password in the adjacent field.

Step 10 Click **Save**.

Step 11 Click **Network**.

Step 12 Click **Add Network** > **Add Object**.

Step 13 In the **Name** field, enter a name to identify the object (for example, **`MyCaptivePortalNetwork`**).

Step 14 Click **FDQN** and, in the field, enter the name of the captive portal's FDQN.

Step 15 Click an option for **Lookup**.

The following figure shows an example.

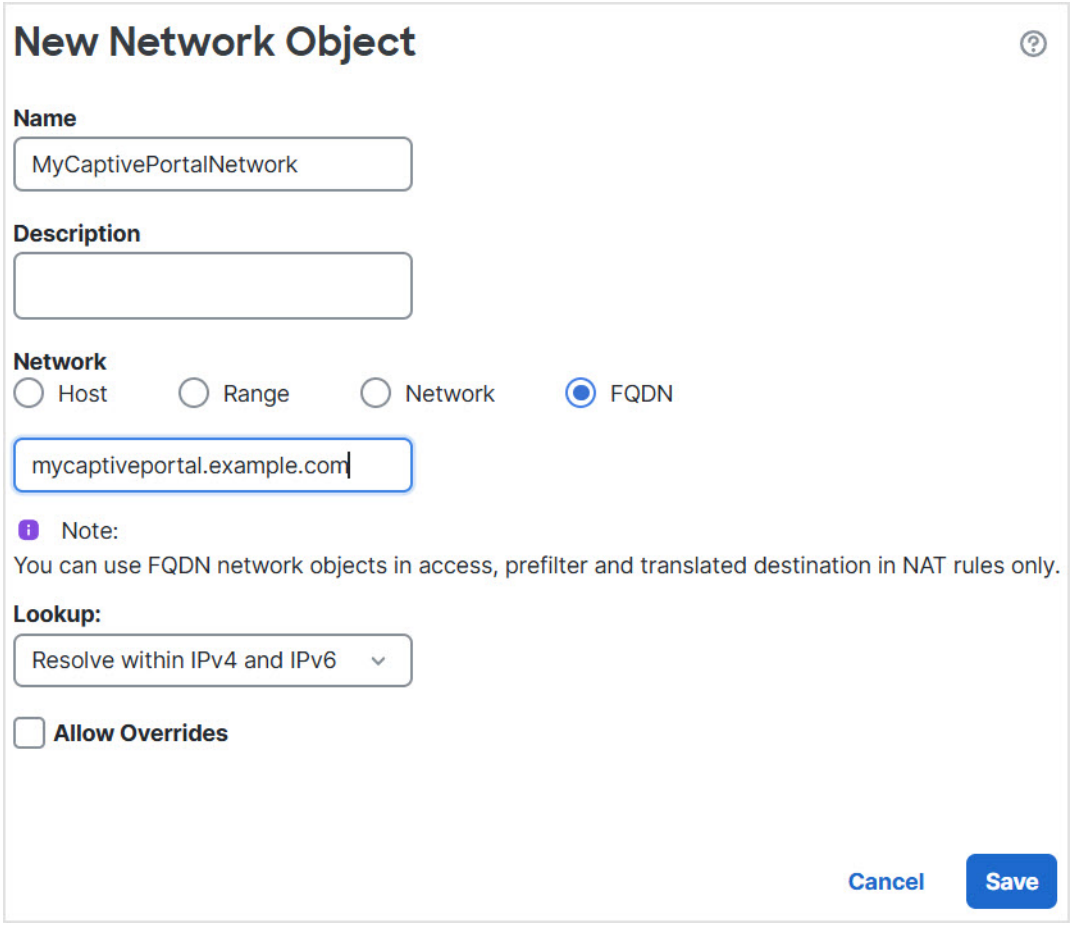

Step 16 Click **Save**.

What to do next

Configure the Captive Portal Part 2: Create an Identity Policy and Active Authentication Rule, on page 2532

Configure the Captive Portal Part 2: Create an Identity Policy and Active Authentication Rule

Before you begin

This multi-part procedure shows how to set up the captive portal using the default TCP port 885 and using a Firewall Management Center server certificate for both the captive portal and for TLS/SSL decryption. Each part of this example explains one task required to enable the captive portal to perform active authentication.

If you follow all the steps in this procedure, you can configure captive portal to work for users in your domains. You can optionally perform additional tasks, which are discussed in each part of the procedure.

For an overview of the entire procedure, see How to Configure the Captive Portal for User Control, on page 2529.

Procedure

Step 1 Log in to the Firewall Management Center if you have not already done so.

Step 2 Click **Policies** > **Access Control heading** > **Identity** and create or edit an identity policy.

Step 3 (Optional.) Click **Add Category** to add a category for the captive portal identity rules and enter a **Name** for the category.

Step 4 Click the **Active Authentication** tab.

Step 5 Choose the appropriate *root* **Server Certificate** from the list or click **Add** (+) to add a certificate.

For more information, see Connect Securely to Active Directory or LDAP, on page 2442.

Note
Captive portal does *not* support the use of Digital Signature Algorithm (DSA) or Elliptic Curve Digital Signature Algorithm (ECDSA) certificates.

Step 6 From the **Redirect to Host Name** field, click the network object you previously created or click **Add** (+).

Step 7 Enter **885** in the **Port** field and specify the **Maximum login attempts**.

Step 8 Uncheck **Share active authentication across firewalls** to enable the Firewall Management Center to require users to reauthenticate every time they access your network using a different managed device than the last time.
For more information about this option, see Captive Portal Fields, on page 2539.

Step 9 (Optional.) Choose an **Active Authentication Response Page** as described in Captive Portal Fields, on page 2539.

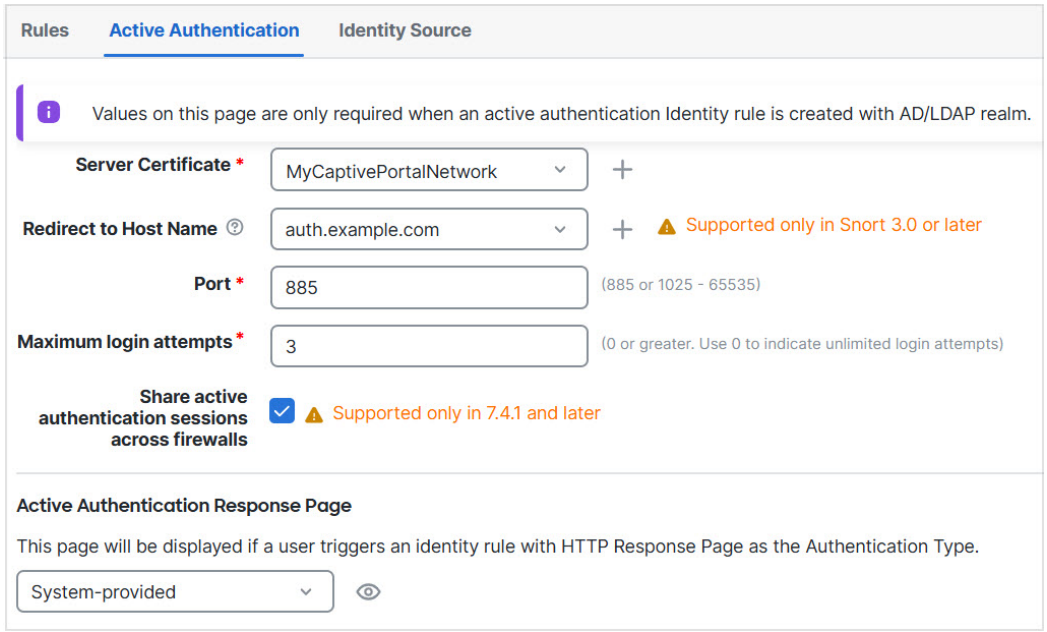

Step 10 (If you upgraded to version 7.4.1 from an earlier version only and you authenticate users with a realm sequence.) Click **Edit** (✎) and see Update a Custom Authentication Form, on page 2534.

Step 11 Click **Save**.

Step 12 Click **Rules**.

Step 13 Click **Add Rule** to add a new captive portal identity policy rule, or click **Edit** (✎) to edit an existing rule.

Step 14 Enter a **Name** for the rule.

Step 15 From the **Action** list, choose **Active Authentication**.

Step 16 Click **Realm & Settings**.

Step 17 From the **Realms** list, choose a realm or realm sequence to use for user authentication.

Step 18 (Optional.) Check **Identify as Guest if authentication cannot identify user**. For more information, see Captive Portal Fields, on page 2539.

Step 19 Choose an **Authentication Protocol** from the list.

You *cannot* authenticate users with a realm sequence if you choose **NTLM**, **Kerberos**, or **HTTP Negotiate** authentication protocols. Choose **HTTP Basic** or **HTTP Response Page** instead.

Step 20 (Optional.) To exempt specific application traffic from captive portal, see Exclude Applications from Captive Portal, on page 2540.

Step 21 Add conditions to the rule (port, network, and so on) as discussed in Identity Rule Conditions, on page 2557.

Step 22 Click **Add**.

Step 23 At the top of the page, click **Save**.

What to do next

Continue with Configure the Captive Portal Part 3: Create a TCP Port Access Control Rule, on page 2534.

Update a Custom Authentication Form

After you upgrade to version 7.4.1 (or later) from an earlier release, you must add the following line to a custom authentication form for users to see a list of domains when they authenticate with captive portal. (This task is always required if you use the HTTP Response Page authentication type; if users authenticate with a realm using another authentication type, this task is optional.)

On the **Active Authentication** tab page of your identity rule, click **Edit** (✎) and enter the following in the part of the form that requires the user to log in:

```
<select name="realm" id="realm"></select>
```

Configure the Captive Portal Part 3: Create a TCP Port Access Control Rule

This part of the procedure shows how to create an access control rule that allows the captive portal to communicate with clients using TCP port 885, which is the captive portal's default port. You can choose another port if you wish, but the port must match the one you chose in Configure the Captive Portal Part 2: Create an Identity Policy and Active Authentication Rule, on page 2532.

Before you begin

For an overview of the entire captive portal configuration, see How to Configure the Captive Portal for User Control, on page 2529.

Procedure

Step 1 Log in to the Firewall Management Center if you have not already done so.

Step 2 If you haven't done so already, create a certificate for the captive portal as discussed in PKI, on page 1404.

Step 3 Click **Policies** > **Access Control heading** > **Access Control** and create or edit an access control policy.

Step 4 Click **Add Rule**.

Step 5 Enter a **Name** for the rule.

Step 6 Choose **Allow** from the **Action** list.

Step 7 Click **Ports**.

Step 8 From the **Protocol** list under the **Selected Destination Ports** field, choose **TCP**.

Step 9 In the **Port** field, enter **885**.

Step 10 Click **Add** next to the **Port** field.
The following figure shows an example.

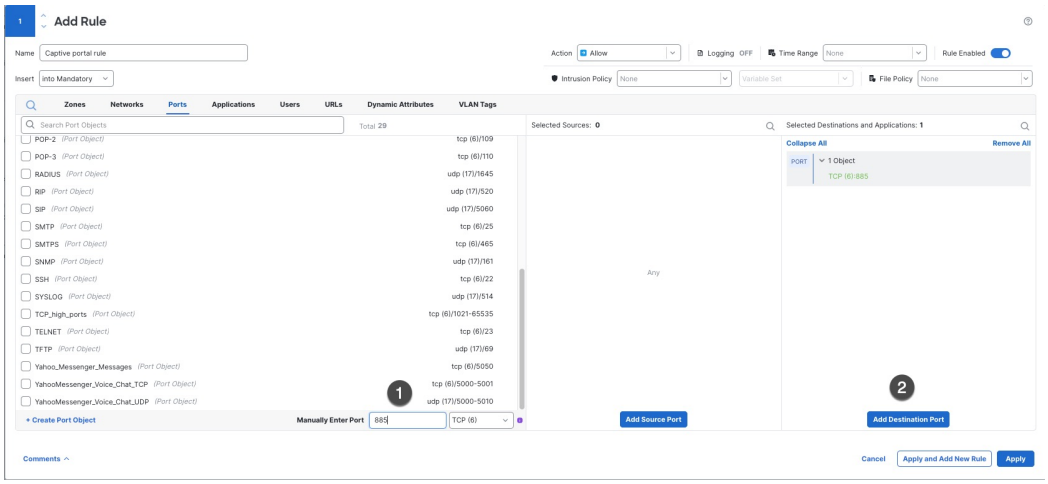

Step 11 Click **Add** at the bottom of the page.

What to do next

Continue with Configure the Captive Portal Part 4: Create a User Access Control Rule, on page 2535.

Configure the Captive Portal Part 4: Create a User Access Control Rule

This part of the procedure discusses how to add an access control rule that enables users in a realm to authenticate using captive portal.

Before you begin

For an overview of the entire captive portal configuration, see How to Configure the Captive Portal for User Control, on page 2529.

Procedure

Step 1 In the rule editor, click **Add Rule**.

Step 2 Enter a **Name** for the rule.

Step 3 Choose **Allow** from the **Action** list.

Step 4 Click **Users**.

Step 5 In the **Available Realms** list, click the realms to allow.

Step 6 If no realms display, click **Refresh** ( ).

Step 7 In the **Available Users** list, choose the users to add to the rule and click **Add to Rule**.

Step 8 (Optional.) Add conditions to the access control policy as discussed in Identity Rule Conditions, on page 2557.

Step 9 Click **Add**.

Step 10 On the access control rule page, click **Save**.

Step 11 In the policy editor, set the rule position. Click and drag or use the right-click menu to cut and paste. Rules are numbered starting at 1. The system matches traffic to rules in top-down order by ascending rule number. The first rule that traffic matches is the rule that handles that traffic. Proper rule order reduces the resources required to process network traffic and prevents rule preemption.

What to do next

Captive Portal Example: Create a Decryption Policy with an Outbound Rule, on page 2536

Captive Portal Example: Create a Decryption Policy with an Outbound Rule

This part of the procedure discusses how to create a decryption policy to decrypt and resign traffic before the traffic reaches the captive portal. The captive portal can authenticate traffic only after it has been decrypted.

Before you begin

You must have an internal certificate authority (CA) for your outbound server; in other words, the managed device that decrypts the traffic for captive portal users to authenticate. This certificate must be different from the *internal certificate* you use to authenticate the captive portal with the managed device.

Procedure

Step 1 Log in to the Firewall Management Center if you haven't already done so.

Step 2 Click **Policies** > **Access Control heading** > **Decryption**.

Step 3 Click **New Policy**.

Step 4 Give the policy a unique **Name** and, optionally, a **Description**.

Step 5 Click the **Outbound Connections** tab.

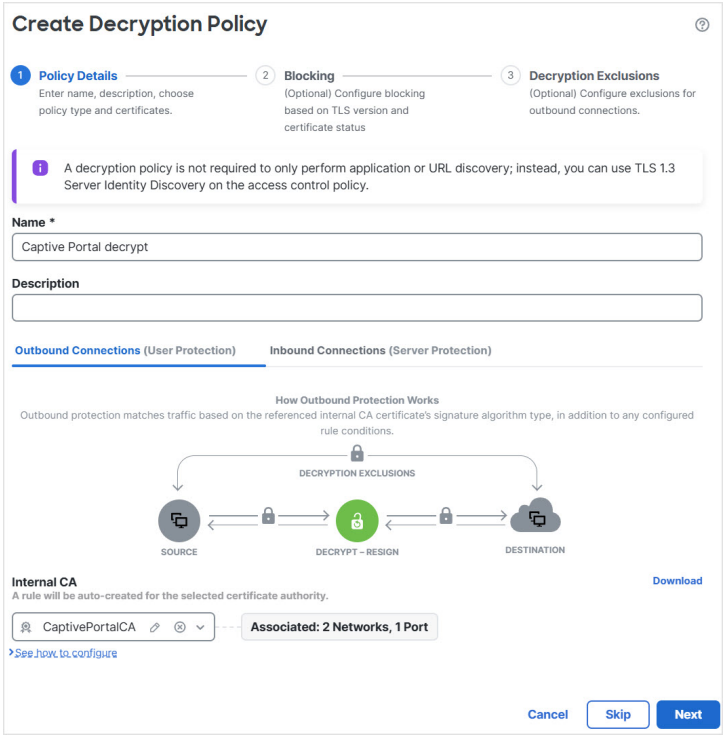

Step 6 Upload or choose certificates for the rules.

The system creates one rule per combination of CA and networks/ports.

Step 7 (Optional.) Choose networks and ports.

For more information:

- Decryption Rule Conditions, on page 2316
- Network Rule Conditions, on page 939
- Port Rule Conditions, on page 940

Step 8 Click **Save**.

Step 9 Click **Edit** (✎) next to the decryption policy you just created.

Step 10 Click **Edit** (✎) next to the decryption rule for captive portal.

Step 11 Click **Users**.

Step 12 Above the **Available Realms** list, click **Refresh** (↻).

Step 13 In the **Available Realms** list, click **Special Identities**.

Step 14 In the **Available Users** list, click **Unknown**.

Step 15 Click **Add to Rule**.

The following figure shows an example.

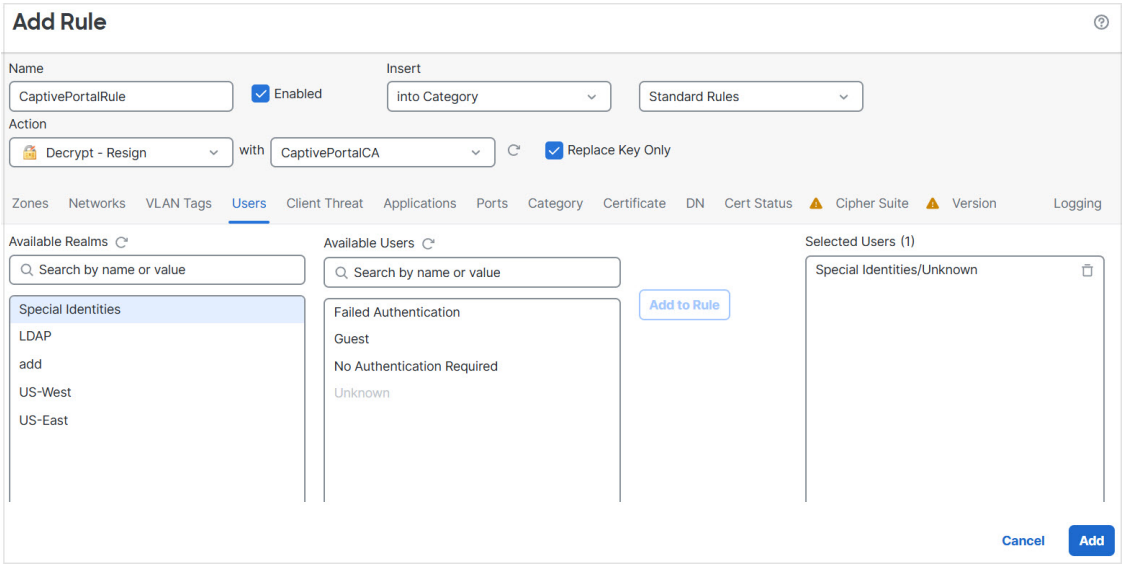

Step 16 (Optional.) Set other options as discussed in Decryption Rule Conditions, on page 2316.

Step 17 Click **Add**.

What to do next

Configure Captive Portal Part 6: Associate Identity and Decryption Policies with the Access Control Policy, on page 2538

Configure Captive Portal Part 6: Associate Identity and Decryption Policies with the Access Control Policy

This part of the procedure discusses how to associate the identity policy and TLS/SSL **Decrypt - Resign** rule with the access control policy you created earlier. After this, users can authenticate using the captive portal.

Before you begin

For an overview of the entire captive portal configuration, see How to Configure the Captive Portal for User Control, on page 2529.

Procedure

Step 1 Click **Policies** > **Access Control heading** > **Access Control** and edit the access control policy you created as discussed in Configure the Captive Portal Part 3: Create a TCP Port Access Control Rule, on page 2534. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 2 Either create a new access control policy or edit an existing policy.

Step 3 At the top of the page, click the word **Identity**.

Step 4 From the list, choose the name of your identity policy and, at the top of the page, click **Save**.

Step 5 Repeat the preceding steps to associate your captive portal decryption policy with the access control policy.

Step 6 If you haven't done so already, target the policy at managed devices as discussed in Assigning Devices to an Access Control Policy, on page 1755.

What to do next

- Deploy your identity and access control policies to managed devices as discussed in Deploy Configuration Changes, on page 259.
- Monitor user activity as discussed in *Using Workflows* in the Cisco Secure Firewall Management Center Administration Guide.

Captive Portal Fields

Use the following fields to configure captive portal on the **Active Authentication** tab page of your identity policy. See also Identity Rule Fields, on page 2564 and Exclude Applications from Captive Portal, on page 2540.

Server Certificate

An internal certificate presented by the captive portal daemon.

Note Captive portal does *not* support the use of Digital Signature Algorithm (DSA) or Elliptic Curve Digital Signature Algorithm (ECDSA) certificates.

Port

The port number to use for the captive portal connection. You must set up your access control rule with a TCP port to use for the captive portal, then associate the identity policy with that access control policy. For more information, see Configure the Captive Portal Part 3: Create a TCP Port Access Control Rule, on page 2534.

Maximum login attempts

The maximum allowed number of failed login attempts before the system denies a user's login request.

Share active authentication sessions across firewalls

Determines whether or not users are required to authenticate when their authentication session is sent to a different managed device than one they previously connected to. If your organization requires users to authenticate every time they change locations or sites, you should *disable* this option.

- (Default, which continues previous behavior.) Check the box to allow users to authenticate with any managed device associated with the active authentication identity rule.
- Uncheck the box to require the user to authenticate with a different managed device, even if they have already authenticated with another managed device to which the active authentication rule is deployed. Use this option if your organization requires authentication per location or site and a managed device is deployed per site.

Managed devices are either clustered or are in a high availability pair act as if they are the same device; in particular,

- Managed devices in the same cluster or high availability pair: Saves the user session to maintain consistency across the pair. On failover, the secondary has the current user session data.
- Managed devices in a different cluster or high availability pair: User session data is not shared with and therefore not stored on these devices.

Active Authentication Response Page

The system-provided or custom HTTP response page you want to display to captive portal users. After you select an **Active Authentication Response Page** in your identity policy active authentication settings, you also must configure one or more identity rules with **`HTTP Response Page`** as the **Authentication Protocol**.

The system-provided HTTP response page includes **Username** and **Password** fields, a list of realms if you chose to authenticate with a realm sequence, as well as a **Login as guest** button to allow users to access the network as guests. To display a single login method, configure a custom HTTP response page.

An example of what users see when logging in with a response page is shown in Create a Sample Identity Policy with an Active Authentication Rule, on page 2568.

Choose the following options:

- To use a generic response, click **System-provided**. You can click **View** ( ) to view the HTML code for this page.
- To create a custom response, click **Custom**. A window with system-provided code is displayed that you can replace or modify. When you are done, save your changes. You can edit a custom page by clicking **Edit** ( ).

Related Topics

Internal Certificate Objects, on page 1412

Exclude Applications from Captive Portal

You can select applications (identified by their HTTP `User-Agent` strings) and exempt them from captive portal active authentication. This allows traffic from the selected applications to pass through the identity policy without authenticating.

Note Only applications with the **`User-Agent Exclusion`** **Tag** are displayed in this list.

Procedure

Step 1 If you haven't done so already, log in to the Firewall Management Center.

Step 2 Click **Policies** > **Access Control heading** > **Identity** .

Step 3 Edit the identity policy that contains the captive portal rule.

Step 4 On **Realm & Settings** tab page, expand **HTTP User Agent Exclusions**.

- In the first column, select the check box next to each item to filter applications, then one or more applications, and click **Add to Rule**.

  Check boxes are ANDed together.

- To narrow the filters that are displayed, type a search string in the **Search by name** field; this is especially useful for categories and tags. To clear the search, click **Clear** ( ).
- To refresh the filters list and clear any selected filters, click **Reload** ( ).

Note
The list displays 100 applications at a time.

Step 5 Choose the applications that you want to add to the filter from the **Available Applications** list:

- To narrow the individual applications that appear, enter a search string in the **Search by name** field. To clear the search, click **Clear** ( ).
- Use paging at the bottom of the list to browse the list of individual available applications.
- To refresh the applications list and clear any selected applications, click **Reload** ( ).

Step 6 Add the selected applications to exclude from external authentication. You can click and drag, or you can click **Add to Rule**. The result is the combination of the application filters you selected.

What to do next

- Continue configuring the identity rule as described in Create an Identity Rule, on page 2563.

Troubleshoot the Captive Portal Identity Source

For other related troubleshooting information, see Troubleshoot Realms and User Downloads, on page 2456 and Troubleshoot User Control, on page 2573.

If you experience issues with captive portal, check the following:

- The time on your captive portal managed device must be synchronized with the time on the Firewall Management Center.
- If you have DNS resolution configured and you create an identity rule to perform **`Kerberos`** (or **`HTTP Negotiate`**, if you want Kerberos as an option) captive portal, you must configure your DNS server to resolve the fully qualified domain name (FQDN) of the captive portal device. The FQDN must match the hostname you provided when configuring DNS.

  For more information, see About Hostname Redirect, on page 2526.

- If you're using Kerberos authentication, the managed device's host name must be less than 15 characters (it's a NetBIOS limitation set by Windows); otherwise, captive portal authentication fails. You set the managed device host name when you set up the device. For more information, see an article like this one on the Microsoft documentation site: Naming conventions in Active Directory for computers, domains, sites, and OUs.

- DNS must return a response of 64KB or less to the hostname; otherwise, the AD connection test fails. This limit applies in both directions and is discussed in RFC 6891 section-6.2.5.
- If the captive portal is configured correctly but the redirect to an IP address or fully-qualified domain name (FQDN) fails, disable endpoint security software. This type of software can interfere with the redirection.
- If you select **`Kerberos`** (or **`HTTP Negotiate`**, if you want Kerberos as an option) as the **Authentication Type** in an identity rule, the **Realm** you select must be configured with an **AD Join Username** and **AD Join Password** to perform Kerberos captive portal active authentication.
- If you select **`HTTP Basic`** as the **Authentication Type** in an identity rule, users on your network might not notice their sessions time out. Most web browsers cache the credentials from **`HTTP Basic`** logins and use the credentials to seamlessly begin a new session after an old session times out.
- If the connection between your Firewall Management Center and a managed device fails, no captive portal logins reported by the device can be identified during the downtime, unless the users were previously seen and downloaded to the Firewall Management Center. The unidentified users are logged as Unknown users on the Firewall Management Center. After the downtime, the Unknown users are reidentified and processed according to the rules in your identity policy.
- If the device you want to use for captive portal contains both inline and routed interfaces, you must configure a zone condition in your captive portal identity rules to target only the routed interfaces on the captive portal device.
- The host name of the managed device must be less than 15 characters for Kerberos authentication to succeed.
- To securely log a user out of captive portal and prevent the user from logging in again, an administrator can close their session: **Analysis** > **Users heading** > **Active Sessions**. If the user closes their browser, the browser does not reauthenticate automatically.
- Active FTP sessions are displayed as the **Unknown** user in events. This is normal because, in active FTP, the server (not the client) initiates the connection and the FTP server should not have an associated user name. For more information about active FTP, see RFC 959.
- When the captive portal authenticates users that match an identity rule, any user in a Microsoft Active Directory or LDAP group that has not been downloaded is identified as Unknown. To avoid users being identified as Unknown, configure the realm or realm sequence to download users in all groups you expect to authenticate with captive portal. Unknown users are handled according to the associated access control policy; if the access control policy is configured to block Unknown users, these users are blocked.

  To make sure the system downloads all users in a realm or realm sequence, make sure the groups are in the Available Groups list in the realm's configuration.

  For more information, see Synchronize Users and Groups, on page 2445.

History for Captive Portal

Feature	Minimum Firewall Management CenterFirewall Management Center	Minimum Firewall Threat Defense	Details
Authentication using a realm or realm sequence.	7.4.1	7.4.1	You can configure active authentication for either an LDAP realm; or a Microsoft Active Directory realm or a realm sequence. In addition, you can configure a passive authentication rule to fall back to active authentication using either a realm or a realm sequence. You can optionally share sessions between managed devices that share the same identity policy in access control rules. In addition, you have the option to require users to authenticate again when they access the system using a different managed device than they accessed previously. Microsoft Azure Active Directory cannot be used with captive portal. New/modified screens: • **Policies** > **Identity** > **(edit policy)** > **Active Authentication** > **Share active authentication sessions across firewalls** • **Identity policy** > **(edit)** > **Add Rule** > **Passive Authentication** > **Realms & Settings** > **Use active authentication if passive or VPN identity cannot be established** • **Identity policy** > **(edit)** > **Add Rule** > **Active Authentication** > **Realms & Settings** > **Use active authentication if passive or VPN identity cannot be established**
Share active authentication sessions across firewalls.	7.4.1	7.4.1	Determines whether or not users are required to authenticate when their authentication session is sent to a different managed device than one they previously connected to. If your organization requires users to authenticate every time they change locations or sites, you should *disable* this option. • (Default.) Enable to allow users to authenticate with any managed device associated with the active authentication identity rule. • Disable to require the user to authenticate with a different managed device, even if they have already authenticated with another managed device to which the active authentication rule is deployed. New/modified screens: **Policies** > **Identity** > **(edit policy)** > **Active Authentication** > **Share active authentication sessions across firewalls**
Hostname redirect.	7.1.0	7.1.0 with Snort 3	You can use a network object that contains the fully-qualified host name (FQDN) of the interface that captive portal can use for active authentication requests.

Feature	Minimum Firewall Management CenterFirewall Management Center	Minimum Firewall Threat Defense	Details
Guest login.	6.1.0	6.1.0	Users can log in as guest using captive portal.
Captive portal.	6.0.0	6.0.0	Feature introduced. You can use the captive portal to require users to enter their credentials when prompted in a browser window. The mapping also allows policies to be based on a user or group of users.

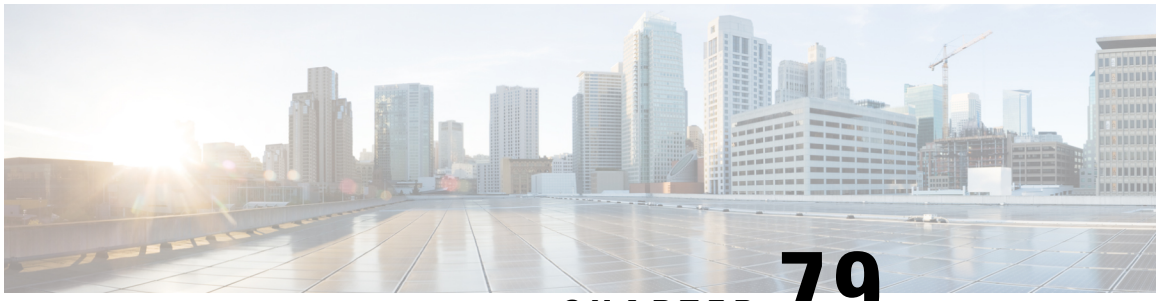

CHAPTER 79

User Control with Remote Access VPN

The following topics discuss how to perform user awareness and user control with Remote Access VPN:

- The Remote Access VPN Identity Source, on page 2545
- Configure RA VPN for User Control, on page 2546
- Troubleshoot the Remote Access VPN Identity Source, on page 2546
- History for RA VPN, on page 2548

The Remote Access VPN Identity Source

Secure Client is the only client supported on endpoint devices for remote VPN connectivity to Firewall Threat Defense devices.

When you set up a secure VPN gateway as discussed in Create a New Remote Access VPN Policy, on page 1584, you can set up an identity policy for those users and associate the identity policy with an access control policy, provided your users are in an Active Directory repository.

Note If you use remote access VPN with User Identity and RADIUS as the identity source, you must configure the realm (**Objects** > **Object Management** > **AAA Server** > **RADIUS Server Group**).

The login information provided by a remote user is validated by an LDAP or AD realm or a RADIUS server group. These entities are integrated with the Secure Firewall Threat Defense secure gateway.

Note If users authenticate with remote access VPN using Active Directory as the authentication source, users must log in using their username; the format `domain\username` or `username@domain` fails. (Active Directory refers to this username as the *logon name* or sometimes as `sAMAccountName`.) For more information, see User Naming Attributes on MSDN.

If you use RADIUS to authenticate, users can log in with any of the preceding formats.

Once authenticated via a VPN connection, the remote user takes on a *VPN Identity*. This VPN Identity is used by *identity policies* on the Secure Firewall Threat Defense secure gateway to recognize and filter network traffic belonging to that remote user.

Identity policies are associated with access control policies, which determine who has access to network resources. It is in this way that the remote user blocked or allowed to access your network resources.

Related Topics

VPN Overview, on page 1497
Remote Access VPN Overview, on page 1573
VPN Basics, on page 1498
Remote Access VPN Features, on page 1574
Guidelines and Limitations for Remote Access VPNs, on page 1580
Create a New Remote Access VPN Policy, on page 1584

Configure RA VPN for User Control

Before you begin

- Create a realm as discussed in Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.
- To use authentication, authorization, and auditing (AAA), set up a RADIUS server group as discussed in Add a RADIUS Server Group, on page 1364.

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Devices** > **VPN** > **Remote Access**.

Step 3 See Create a New Remote Access VPN Policy, on page 1584.

What to do next

- Specify users to control and other options using an identity policy as described in Create an Identity Policy, on page 2555.
- Associate the identity rule with an access control policy, which filters and optionally inspects traffic, as discussed in Associating Other Policies with Access Control, on page 1762.
- Deploy your identity and access control policies to managed devices as discussed in Deploy Configuration Changes, on page 259.
- Monitor VPN user traffic as discussed in VPN Session and User Information, on page 1710.

Troubleshoot the Remote Access VPN Identity Source

- For other related troubleshooting information, see Troubleshoot Realms and User Downloads, on page 2456 and Troubleshoot User Control, on page 2573.

- If you experience issues with Remote Access VPN, check the connection between your Firewall Management Center and a managed device. If the connection fails, all Remote Access VPN logins reported by the device cannot be identified during the downtime, unless the users were previously seen and downloaded to the Firewall Management Center.

  The unidentified users are logged as Unknown users on the Firewall Management Center. After the downtime, the Unknown users are re identified and processed according to the rules in your identity policy.

- The host name of the managed device must be less than 15 characters for Kerberos authentication to succeed.

- Active FTP sessions are displayed as the **Unknown** user in events. This is normal because, in active FTP, the server (not the client) initiates the connection and the FTP server should not have an associated user name. For more information about active FTP, see RFC 959.

Not Observing Correct Settings for VPN Statistics

This task discusses steps you must take after either enabling or disabling the **VPN Statistics** setting in a health policy. Failure to perform this task means managed devices have a health policy with incorrect settings.

Procedure

Step 1 Log in to the Secure Firewall Management Center if you haven't already done so.

Step 2 Click **System** (⚙) > **Health** > **Policy**.

Step 3 Under Firewall Threat Defense Health Policies, click **Edit** (✎) next to a policy to edit.

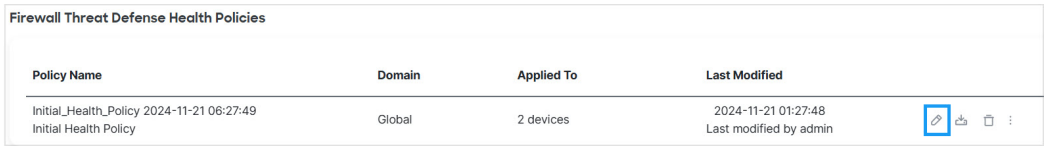

Step 4 On the **Health Modules** tab page, scroll down to locate **VPN Statistics**.

Step 5 Verify the VPN statistics setting is correct or change it if necessary.

Step 6 If you changed the setting, click **Save**, then click **Cancel** to return to the health policy.

Step 7 Under Firewall Threat Defense Health Policies, click **Deploy health policy** (⬆) to apply the policy.

Step 8 In the **Policy Assignments & Deploy** dialog box, move the devices to which to deploy the health policy to the **Selected Devices** field.

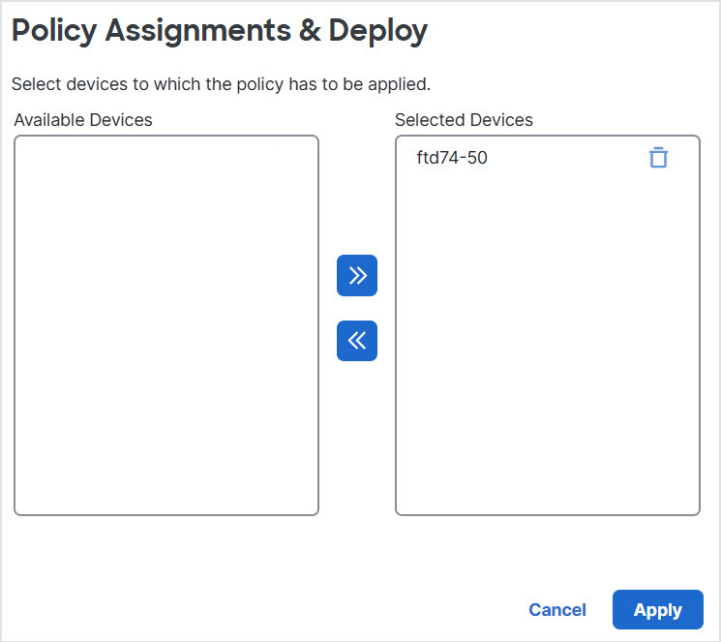

Step 9 Click **Apply**.
A message is displayed when the health policy is deployed.

Step 10 After the health policy has finished deploying, click **Policies** > **Access Control heading** > **Access Control** to edit an access control policy.

Step 11 Click **Edit** (✎) next to a policy to edit.

Step 12 Make a minor change to the policy, such as changing its name.

Step 13 Save the access control policy.

Step 14 Deploy configuration changes; see Deploy Configuration Changes, on page 259..

History for RA VPN

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Remote Access VPN	6.2.1	Any	Feature introduced. RA VPN allows individual users to connect to a private business network from a remote location using a laptop or desktop computer connected to the internet, or an Android or Apple iOS mobile device. Remote users transfer data securely and confidentially using encryption techniques crucial for data being transferred over shared mediums and the Internet.

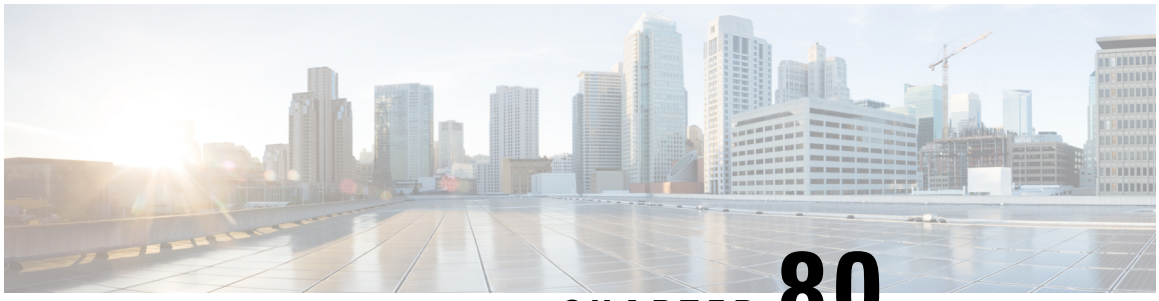

CHAPTER 80

User Control with TS Agent

To use the TS Agent as an identity source for user awareness and user control, install and configure the TS Agent software as discussed in the Cisco Terminal Services (TS) Agent Guide.

What to do next:

- Specify users to control and other options using an identity policy as described in Create an Identity Policy, on page 2555.
- Associate the identity rule with an access control policy, which filters and optionally inspects traffic, as discussed in Associating Other Policies with Access Control, on page 1762.
- Deploy your identity and access control policies to managed devices as discussed in Deploy Configuration Changes, on page 259.
- Monitor user activity as discussed in *Using Workflows* in the Cisco Secure Firewall Management Center Administration Guide .

- The Terminal Services (TS) Agent Identity Source, on page 2549
- TS Agent Guidelines, on page 2550
- User Control with TS Agent, on page 2550
- Troubleshoot the TS Agent Identity Source, on page 2550
- History for TS Agent, on page 2551

The Terminal Services (TS) Agent Identity Source

The TS Agent is a passive authentication method and one of the authoritative identity sources supported by the system. A Windows Terminal Server performs the authentication, and the TS Agent reports it to a standalone or high availability Firewall Management Center.

When installed on Windows Terminal Servers, the TS Agent assigns a unique port range to individual users as they log in or log out of a monitored network. The Firewall Management Center uses the unique port to identify individual users in the system. You can use one TS Agent to monitor user activity on one Windows Terminal Server and send encrypted data to a Firewall Management Center.

The TS Agent does not report failed login attempts. The data gained from the TS Agent can be used for user awareness and user control.

TS Agent Guidelines

The TS Agent requires a multi-step configuration, and includes the following:

1. A Windows Terminal Server with the TS Agent installed and configured.
2. One or more identity realms targeting the users your server is monitoring.

You install the TS Agent on a Microsoft Windows Terminal Server. For detailed information about the multi-step TS Agent installation and configuration and a complete discussion of the server and system requirements, see the Cisco Terminal Services (TS) Agent Guide.

TS Agent data is visible in the Users, User Activity, and Connection Event tables and can be used for user awareness and user control.

Note If the TS Agent monitors the same users as another passive authentication identity source (ISE/ISE-PIC), the Firewall Management Center prioritizes the TS Agent data. If the TS Agent and another passive identity source report activity by the same IP address, only the TS Agent data is logged to the Firewall Management Center.

User Control with TS Agent

To use the TS Agent as an identity source for user awareness and user control, install and configure the TS Agent software as discussed in the Cisco Terminal Services (TS) Agent Guide.

What to do next:

- Specify users to control and other options using an identity policy as described in Create an Identity Policy, on page 2555.
- Associate the identity rule with an access control policy, which filters and optionally inspects traffic, as discussed in Associating Other Policies with Access Control, on page 1762.
- Deploy your identity and access control policies to managed devices as discussed in Deploy Configuration Changes, on page 259.
- Monitor user activity as discussed in *Using Workflows* in the Cisco Secure Firewall Management Center Administration Guide .

Troubleshoot the TS Agent Identity Source

For other related troubleshooting information, see Troubleshoot Realms and User Downloads, on page 2456 and Troubleshoot User Control, on page 2573.

If you experience issues with the TS Agent integration, check:

- You must synchronize the time on your TS Agent server with the time on the Firewall Management Center.

- If the TS Agent monitors the same users as another passive authentication identity source (ISE/ISE-PIC), the Firewall Management Center prioritizes the TS Agent data. If the TS Agent and a passive identity source report activity by the same IP address, only the TS Agent data is logged to the Firewall Management Center.
- Active FTP sessions are displayed as the **Unknown** user in events. This is normal because, in active FTP, the server (not the client) initiates the connection and the FTP server should not have an associated user name. For more information about active FTP, see RFC 959.

For more troubleshooting information, see the Cisco Terminal Services (TS) Agent Guide.

History for TS Agent

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
TS Agent for user control.	7.2.0	6.2.0	Feature introduced. Firepower now provides the ability to better identify individual users in shared environments, such as Citrix's Virtual Desktop Infrastructure (VDI), to accurately enforce user-based policy rules on the firewall. Users are identified by ports used. The TS Agent software is updated independently of the Firepower Management Center. For more information, see: • *Cisco Terminal Services (TS) Agent Guide* available on cisco.com • Cisco Firepower Compatibility Guide

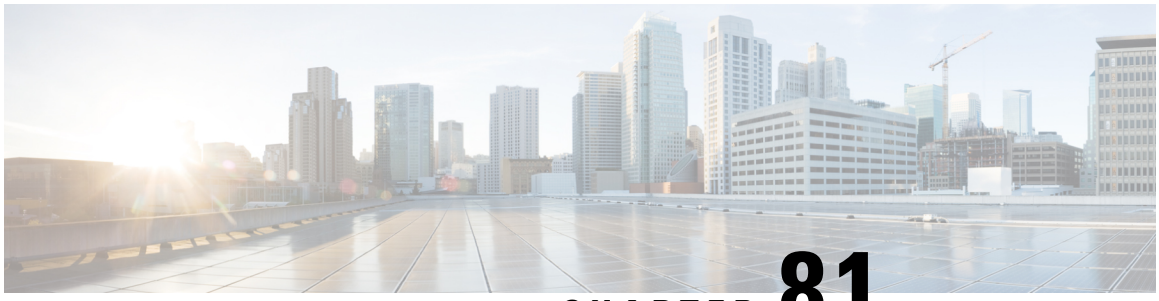

CHAPTER 81

User Identity Policies

The following topics discuss how to create and manage identity rules and identity policies:

- About Identity Policies, on page 2553
- License Requirements for Identity Policies, on page 2554
- Requirements and Prerequisites for Identity Policies, on page 2554
- Create an Identity Policy, on page 2555
- Identity Rule Conditions, on page 2557
- Create an Identity Rule, on page 2563
- Sample Identity Policies and Rules, on page 2565
- Manage an Identity Policy, on page 2572
- Manage an Identity Rule, on page 2573
- Troubleshoot User Control, on page 2573

About Identity Policies

Identity policies contain identity rules. Identity rules associate sets of traffic with a realm and an authentication method: passive authentication, active authentication, or no authentication.

With the exception noted in the following paragraphs, you must configure realms and authentication methods you plan to use before you can invoke them in your identity rules:

- You configure realms outside of your identity policy, at **Integration** > **Other Integrations** > **Realms** > **Realms**. For more information, see Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.
- You configure ISE/ISE-PIC, a passive authentication identity source, at **Integration** > **Other Integrations** > **Identity Sources**. For more information, see Ways to Configure the Cisco Identity Services Engine (Cisco ISE) Identity Source, on page 2512.
- You configure the TS Agent, a passive authentication identity source, outside the system. For more information, see the *Cisco Terminal Services (TS) Agent Guide*.
- You configure captive portal, an active authentication identity source, within the identity policy. For more information, see How to Configure the Captive Portal for User Control, on page 2529.
- You configure Remote Access VPN, an active authentication identity source, in Remote Access VPN policies. For more information, see Remote Access VPN Authentication, on page 1576.

After you add multiple identity rules to a single identity policy, order the rules. The system matches traffic to rules in top-down order by ascending rule number. The first rule that traffic matches is the rule that handles the traffic.

You can also filter traffic by network object, which limits the network each device monitors in the event your devices are at or near their memory limits.

After you configure one or more identity policies, you must associate one identity policy with your access control policy. When traffic on your network matches the conditions in your identity rule, the system associates the traffic with the specified realm and authenticates the users in the traffic using the specified identity source.

If you do not configure an identity policy, the system does not perform user authentication.

Exception to creating an identity policy

An identity policy is not required if all of the following are true:

- You use the ISE/ISE-PIC identity source.
- You do not use users or groups in access control policies.
- You use Security Group Tags (SGT) in access control policies. For more information, see ISE SGT vs Custom SGT Rule Conditions.

Related Topics

How to Set Up an Identity Policy, on page 2378

License Requirements for Identity Policies

Threat Defense License

Any

Requirements and Prerequisites for Identity Policies

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Create an Identity Policy

This task discusses how to create an identity policy.

Before you begin

An identity policy is required to use users and groups in a realm in access control policies. Create and enable one or more realms as described in Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.

(Optional.) If a particular managed device monitors a large number of user groups, the system might drop user mappings based on groups due to managed device memory limitations. As a result, rules with realm or user conditions might not perform as expected. Provided the devices run version 6.7 or later, you can configure the identity rule to monitor traffic by one network or network group object only. To create a network object, see Creating Network Objects, on page 1403.

An identity policy is not required if all of the following are true:

- You use the ISE/ISE-PIC identity source.
- You do not use users or groups in access control policies.
- You use Security Group Tags (SGT) in access control policies. For more information, see ISE SGT vs Custom SGT Rule Conditions.

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Policies** > **Access Control heading** > **Identity** and click **New Policy**.

Step 3 Enter a **Name** and, optionally, a **Description**.

Step 4 Click **Save**.

Step 5 To add a rule to the policy, click **Add Rule** as described in Create an Identity Rule, on page 2563.

Step 6 To create a rule category, click **Add Category**.

Step 7 To configure captive portal active authentication, click **Active Authentication** and see Configure the Captive Portal Part 2: Create an Identity Policy and Active Authentication Rule, on page 2532.

Step 8 (Optional.) To filter traffic by network object, click the **Identity Source** tab. From the list, click the network object to use to filter traffic for this identity policy. Click **Add** (+) to create a new network object.

Step 9 Click **Save** to save the identity policy.

What to do next

- Add rules to your identity policy that specify which users to match and other options; see Create an Identity Rule, on page 2563.
- Associate the identity policy with an access control policy to allow or block selected users from accessing specified resources; see Associating Other Policies with Access Control, on page 1762.

- (Not required for Microsoft Azure AD realms.) Deploy configuration changes to managed devices; see Deploy Configuration Changes, on page 259.

If you encounter issues, see Troubleshoot User Control, on page 2573.

Related Topics

Configure the Captive Portal Part 2: Create an Identity Policy and Active Authentication Rule, on page 2532
Captive Portal Fields, on page 2539
Troubleshoot User Control, on page 2573
Create an Identity Mapping Filter, on page 2556

Create an Identity Mapping Filter

An identity mapping filter can be used to limit the networks to which an identity rule applies. For example, if your Firewall Management Center manages FTDs that have a limited amount of memory, you can limit the networks they monitor.

You must create separate identity mapping filters for IPv4 and IPv6 addresses.

You can also optionally exclude subnets from the following:

- Receiving user-to-IP and Security Group Tag (SGT)-to-IP mappings from ISE.
- Receiving user-to-IP mappings from the passive identity agent.

You should typically do this for lower-memory managed devices to prevent Snort identity health monitor memory errors.

Before you begin

Perform the following tasks:

1. Create a realm, which is required for an identity policy. See Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.
2. Create an identity policy. See Create an Identity Policy, on page 2555.
3. Create a network object or network group object as discussed in Creating Network Objects, on page 1403. The network object or group you create should define the network you want managed devices to monitor in identity policies.

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Policies** > **Access Control heading** > **Identity** .

Step 3 Click **Edit** ( ).

Step 4 Click the **Identity Source** tab.

Step 5 From the **Identity Mapping Filter** list, choose the name of a network object to use as a filter .

To create a new network object, see Creating Network Objects, on page 1403.

Note
To restrict traffic to IPv6 addresses, you must add at least one address, network, or group to the filter.

Step 6 Click **Save**.

Step 7 (Not required for Microsoft Azure AD realms.) Deploy configuration changes to managed devices; see Deploy Configuration Changes, on page 259.

What to do next

(Not required for Microsoft Azure AD realms.) Associate the identity policy with an access control policy as discussed in Associating Other Policies with Access Control, on page 1762.

To check or change ISE identity mapping filters (also referred to as *subnet filters*), use the following commands:

```
show identity-subnet-filter
configure identity-subnet-filter { add | remove } subnet
```

Identity Rule Conditions

Rule conditions enable you to fine-tune your identity policy to target the users and networks you want to control. See one of the following sections for more information.

Related Topics

Security Zone Rule Conditions, on page 1905
Network Rule Conditions, on page 939
VLAN Tags Rule Conditions, on page 1787
Port Rule Conditions, on page 940
Realm & Settings Rule Conditions, on page 2561

Security Zone Rule Conditions

Security zones segment your network to help you manage, classify, and decrypt traffic flow by grouping interfaces across multiple devices.

Security zones control or decrypt traffic by its source and destination security zones. If you add both source and destination zones to a zone condition, matching traffic must originate from an interface in one of the source zones and leave through an interface in one of the destination zones.

Just as all interfaces in a zone must be of the same type (all inline, passive, switched, or routed), all zones used in a zone condition must be of the same type. Because devices deployed passively do not transmit traffic, you cannot use a zone with passive interfaces as a destination zone.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Tip Constraining rules by zone is one of the best ways to improve system performance. If a rule does not apply to traffic through any of device's interfaces, that rule does not affect that device's performance.

Security Zone Conditions and Multitenancy

In a multidomain deployment, a zone created in an ancestor domain can contain interfaces that reside on devices in different domains. When you configure a zone condition in an descendant domain, your configurations apply to only the interfaces you can see.

Network Rule Conditions

Networks control or decrypt traffic by its source and destination IP address, using inner headers. Tunnel rules, which use outer headers, have tunnel endpoint conditions instead of network conditions.

You can use predefined objects to build network conditions, or manually specify individual IP addresses or address blocks.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Note You *cannot* use FDQN network objects in identity rules.

Redirect to Host Name Network Rule Conditions

(Snort 3.0 only.)—You can use a network object that contains the fully-qualified host name (FQDN) of the interface that captive portal can use for active authentication requests.

The FQDN must resolve to the IP address of one of the interfaces on a managed device. By using an FQDN, you can assign a certificate for active authentication that the client will recognize, thus avoiding the untrusted certificate warning users get when being redirected to a managed device's IP address.

The certificate can specify one FQDN, a wildcard FQDN, or multiple FQDNs in the Subject Alternate Names (SAN) in the certificate.

If an identity rule requires active authentication for a user, but you do not specify a redirect FQDN, the user is redirected to the captive portal port on the managed device interface to which they are connected.

If you do not supply a Redirect to Host Name FQDN, the HTTP Basic, HTTP Response Page, and NTLM authentication methods redirect the user to the captive portal using the IP address of the interface. However, for HTTP Negotiate, the user is redirected using the fully-qualified DNS name *firewall-hostname.directory-server-domain-name*. To use HTTP Negotiate without a Redirect to Host Name FQDN, you must also update your DNS server to map this name to the IP addresses of all inside interfaces where you are requiring active authentication. Otherwise, the redirection cannot complete, and users cannot authenticate.

We recommend that you always provide a Redirect to Host Name FQDN to ensure consistent behavior regardless of authentication method.

VLAN Tags Rule Conditions

Note VLAN tags in access rules only apply to inline sets. Access rules with VLAN tags do not match traffic on firewall interfaces.

VLAN rule conditions control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic. The system uses the innermost VLAN tag to filter VLAN traffic, with the exception of the prefilter policy, which uses the outermost VLAN tag in its rules.

Note the following Q-in-Q support:

- Firewall Threat Defense on Firepower 4100/9300—Does not support Q-in-Q (supports only one VLAN tag).
- Firewall Threat Defense on all other models:
  - Inline sets and passive interfaces—Supports Q-in-Q, up to 2 VLAN tags.
  - Firewall interfaces—Does not support Q-in-Q (supports only one VLAN tag).

You can use predefined objects to build VLAN conditions, or manually enter any VLAN tag from 1 to 4094. Use a hyphen to specify a range of VLAN tags.

In a cluster, if you encounter problems with VLAN matching, edit the access control policy advanced options, Transport/Network Preprocessor Settings, and select the **Ignore the VLAN header when tracking connections** option.

Port Rule Conditions

Port conditions allow you to control traffic by its source and destination ports.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Best Practices for Port-Based Rules

Specifying ports is the traditional way to target applications. However, applications can be configured to use unique ports to bypass access control blocks. Thus, whenever possible, use application filtering criteria rather than port criteria to target traffic.

Application filtering is also recommended for applications, like Firewall Threat Defense, that open separate channels dynamically for control vs. data flow. Using port-based access control rules can prevent these kinds of applications from performing correctly, and could result in blocking desirable connections.

Using Source and Destination Port Constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as source port conditions in a single access control rule.

Port, Protocol, and ICMP Code Rule Conditions

Port conditions match traffic based on the source and destination ports. Depending on the rule type, "port" can represent any of the following:

- TCP and UDP—You can control TCP and UDP traffic based on the port. The system represents this configuration using the protocol number in parentheses, plus an optional associated port or port range. For example: TCP(6)/22.
- ICMP—You can control ICMP and ICMPv6 (IPv6-ICMP) traffic based on its internet layer protocol plus an optional type and code. For example: ICMP(1):3:3.
- Protocol—You can control traffic using other protocols that do not use ports.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Best Practices for Port-Based Rules

Specifying ports is the traditional way to target applications. However, applications can be configured to use unique ports to bypass access control blocks. Thus, whenever possible, use application filtering criteria rather than port criteria to target traffic. Note that application filtering is not available in prefilter rules.

Application filtering is also recommended for applications, like FTP, that open separate channels dynamically for control vs. data flow. Using port-based access control rules can prevent these kinds of applications from performing correctly, and could result in blocking desirable connections.

Using Source and Destination Port Constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as destination port conditions in a single access control rule.

Matching Non-TCP Traffic with Port Conditions

You can match non-port-based protocols. By default, if you do not specify a port condition, you are matching IP traffic. Although you can configure port conditions to match non-TCP traffic, there are some restrictions:

- Access control rules—For Classic devices, you can match GRE-encapsulated traffic with an access control rule by using the GRE (47) protocol as a destination port condition. To a GRE-constrained rule, you can add only network-based conditions: zone, IP address, port, and VLAN tag. Also, the system uses outer headers to match **all** traffic in access control policies with GRE-constrained rules. For Firewall Threat Defense devices, use tunnel rules in the prefilter policy to control GRE-encapsulated traffic.
- Decryption rules—These rules support TCP port conditions only.

- IMCP echo—A destination ICMP port with the type set to 0 or a destination ICMPv6 port with the type set to 129 only matches unsolicited echo replies. ICMP echo replies sent in response to ICMP echo requests are ignored. For a rule to match on any ICMP echo, use ICMP type 8 or ICMPv6 type 128.

Realm & Settings Rule Conditions

The **Realm & Settings** tab page enables you to choose a realm or realm sequence to which to apply the identity rule. If you are using captive portal, you have additional options.

Authentication Realm

From the **Realm** list, click a realm or realm sequence.

The realm or realm sequence containing the users you want to perform the specified **Action** on. You must fully configure a realm or realm sequence before selecting it as the realm in an identity rule.

Note If remote access VPN is enabled and your deployment is using a RADIUS server group for VPN authentication, make sure you specify the realm associated with this RADIUS server group.

Active authentication only: other options

If you either choose **Active Authentication** as the authentication type or if you check the box, **Use active authentication if passive or VPN identity cannot be established**, you have the following options.

Use active authentication if passive or VPN identity cannot be established

(Passive authentication rule only.) Selecting this option authenticates users using captive portal active authentication if a passive or a VPN authentication fails to identify them. You must configure an Active Authentication rule in your identity policy in order to select this option. (That is, users must authenticate using the captive portal.)

If you disable this option, users that do not have a VPN identity or that passive authentication cannot identify are identified as Unknown.

Also see the discussion of the **Authentication Realm** list later in this topic,

Identify as Special Identities/Guest if authentication cannot identify user

Selecting this option allows users who fail captive portal active authentication the specified number of times to access your network as a guest. These users appear in the Firewall Management Center identified by their username (if their username exists on the AD or LDAP server) or by **Guest** (if their user name is unknown). Their realm is the realm specified in the identity rule. (By default, the number of failed logins is 3.)

This field is displayed only if you configure **Active Authentication** (that is, captive portal authentication) as the rule **Action**.

Authentication Protocol

The method to use to perform captive portal active authentication. An example of what users see when logging in with a response page is shown in Create a Sample Identity Policy with an Active Authentication Rule, on page 2568.

The selections vary depending on the type of realm, LDAP or AD:

- Choose **HTTP Basic** if you want to authenticate users using an unencrypted HTTP Basic Authentication (BA) connection. Users log in to the network using their browser's default authentication pop-up window.

  Most web browsers cache the credentials from **HTTP Basic** logins and use the credentials to seamlessly begin a new session after an old session times out.

- Choose **NTLM** to authenticate users using a NT LAN Manager (NTLM) connection. This selection is available only when you select an AD realm. If transparent authentication is configured in a user's browser, the user is automatically logged in. If transparent authentication is not configured, users log in to the network using their browser's default authentication pop-up window.
- Choose **Kerberos** to authenticate users using a Kerberos connection. This selection is available only when you select an AD realm for a server with secure LDAP (LDAPS) enabled. If transparent authentication is configured in a user's browser, the user is automatically logged in. If transparent authentication is not configured, users log in to the network using their browser's default authentication pop-up window.

Note The **Realm** you select must be configured with an **AD Join Username** and **AD Join Password** to perform Kerberos captive portal active authentication.

Note If you are creating an identity rule to perform Kerberos captive portal and you have DNS resolution configured, you must configure your DNS server to resolve the fully qualified domain name (FQDN) of the captive portal device. The FQDN must match the host name you provided when configuring DNS.

For Firewall Threat Defense devices, the FQDN must resolve to the IP address of the routed interface used for captive portal.

- Choose **HTTP Negotiate** to allow the captive portal server to choose between HTTP Basic, Kerberos, or NTLM for the authentication connection. This type is available only when you select an AD realm.

Note The **Realm** you choose must be configured with an **AD Join Username** and **AD Join Password** for **HTTP Negotiate** to choose Kerberos captive portal active authentication.

Note If you are creating an identity rule to perform **HTTP Negotiate** captive portal and you have DNS resolution configured, you must configure your DNS server to resolve the fully qualified domain name (FQDN) of the captive portal device. The FQDN of the device you are using for captive portal must match the hostname you provided when configuring DNS.

- Choose **HTTP Response Page** to enable users to choose a realm to log in to.

You can optionally customize the response page; for example, to conform to company style standards.

Active Authentication Realm

(Passive authentication rule only.) If you clicked **Use active authentication if passive or VPN identity cannot be established**, you must click the name of a realm or realm sequence. The availability of a realm or realm sequence is determined by your choice for Authentication Protocol as follows:

- **HTTP Basic** or **HTTP Response Page** authentication protocol: You can choose either a realm or a realm sequence.
- **NTLM**, **Kerberos**, or **HTTP Negotiate** authentication protocol: You can choose a realm only. You *cannot* choose a realm sequence.

Create an Identity Rule

For details about configuration options for identity rules, see Identity Rule Fields, on page 2564.

Before you begin

You must create and enable a realm or realm sequence.

- Create a Microsoft Active Directory realm and realm directory as discussed in Create an LDAP Realm or an Active Directory Realm and Realm Directory, on page 2426.
- (Microsoft AD realm only.) Download users and groups and enable the realm as discussed in Synchronize Users and Groups, on page 2445.
- Create a Microsoft Azure AD (SAML) realm as discussed in Create a Microsoft Azure AD (SAML) Realm, on page 2392.
- (Optional.) Create a realm sequence as discussed in Create a Realm Sequence, on page 2446.
- Rules are evaluated top-down. For a connection that matches the specified network criteria of a given rule, the user is evaluated against the identity realm specified in the rule. If the user is not part of that realm, they will be marked as unknown, and no further rules in the identity policy will be evaluated. Thus, if you have more than one realm that needs to be evaluated, be sure to use realm sequences instead of a single realm.

Caution Adding the first or removing the last active authentication rule when TLS/SSL decryption is disabled (that is, when the access control policy does not include a decryption policy) restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

Note that an active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive or VPN identity cannot be established** selected.

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Policies** > **Access Control heading** > **Identity** .

Step 3 Click **Edit** (✎) next to the identity policy to which to add the identity rule.

If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Click **Add Rule**.

Step 5 Enter a **Name**.

Step 6 If the Specified rule is applicable, check the check box of **Enabled**.

Step 7 To add the rule to an existing category, indicate where you want to **Insert** the rule. To add a new category, click **Add Category**.

Step 8 Choose a rule **Action** from the list.

Step 9 If you're configuring captive portal, see How to Configure the Captive Portal for User Control, on page 2529.

Step 10 (Optional) To add conditions to the identity rule, see Identity Rule Conditions, on page 2557.

Step 11 Click **Add**.

Step 12 In the policy editor, set the rule position. Click and drag or use the right-click menu to cut and paste. Rules are numbered starting at 1. The system matches traffic to rules in top-down order by ascending rule number. The first rule that traffic matches is the rule that handles that traffic. Proper rule order reduces the resources required to process network traffic and prevents rule preemption.

Step 13 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Identity Rule Fields

Use the following fields to configure identity rules.

Enabled

Enabling this option enables the identity rule in the identity policy. Unselecting this option disables the identity rule.

Action

Specify the type of authentication you want to perform on the users in the specified realm: **Passive Authentication** (default), **Active Authentication**, or **No Authentication**. You must fully configure the authentication method, or *identity source*, before selecting it as the action in an identity rule.

Additionally, if VPN is enabled (configured on at least one managed device), remote access VPN sessions are actively authenticated by VPN. Other sessions use the rule action. This means that, if VPN is enabled, VPN identity determination is performed first for all sessions regardless of the selected action. If a VPN

identity is found on the specified realm, this is the identity source used. No additional captive portal active authentication is done, even if selected.

If the VPN identity source is not found, the process continues according to the specified action. You cannot restrict the identity policy to VPN authentication only because if the VPN identity is not found, the rule is applied according to the selected action.

Caution Adding the first or removing the last active authentication rule when TLS/SSL decryption is disabled (that is, when the access control policy does not include an SSL policy) restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

Note that an active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive or VPN identity cannot be established** selected.

For information about which passive and active authentication methods are supported in your version of the system, see About User Identity Sources, on page 2372.

Sample Identity Policies and Rules

The following sections provide examples of configuring an identity policy with either a passive authentication rule or an active authentication rule. In addition, because you can authenticate users with active authentication using either realms or realm sequences, separate examples are provided.

Active authentication means the user authenticates using the captive portal; the user enters network credentials before being allowed to access allowed resources. (RA-VPN is another type of active authentication but it cannot be used together with captive portal authentication. For more information, see The Remote Access VPN Identity Source, on page 2545.)

Passive authentication covers all other types. Passive authentication includes using Microsoft Active Directory realms, Microsoft Azure Active Directory realms, Cisco Identity Services Engine, and others.

You can use a Microsoft Azure Active Directory realm to authenticate users a different way that is not discussed here. For more information, see Create a Microsoft Azure AD (SAML) Realm, on page 2392.

Assumptions

The examples use the following assumptions:

- Microsoft Active Directory (AD) realm named forest.example.com with two child domains configured in a trust relationship:
  - US-West
  - US-East
- A realm sequence named US that includes both realms
- Passive authentication rule that authenticates users with a realm sequence
- Two active authentication rules:

- One rule that authenticates users with a realm and uses the NTLM authentication protocol
- One rule that authenticates users with a realm sequence and uses the HTTP Response Page authentication protocol

- Each sample identity rule is associated with a different identity policy

Passive authentication identity rules

When you configure a Passive Authentication identity rule, you can choose to authenticate users either with LDAP, a Microsoft Active Directory realm, or a Microsoft AD realm sequence. You can use a realm to authenticate with any authentication type; realm sequence restricts the authentication types you can use. For an example, see Create an Identity Policy with a Passive Authentication Rule, on page 2566.

Active authentication identity rules

When you configure an Active Authentication identity rule, you can choose to authenticate users either with LDAP, a Microsoft Active Directory realm, or a Microsoft AD realm sequence. You can use a realm to authenticate with any authentication type; realm sequence restricts the authentication types you can use.

You can also authenticate users with a Microsoft Active Directory realm sequence except for the following authentication types:

- NTLM
- Kerberos
- HTTP Negotiate

For an example, see Create a Sample Identity Policy with an Active Authentication Rule, on page 2568.

Create an Identity Policy with a Passive Authentication Rule

This task discusses how to create an identity policy with a passive authentication rule that authenticates users using the US realm sequence. If a user is not found in the first realm in the sequence, the system searches the other realms in the sequence in the order listed in the realm sequence. If a user is still not found in the realm or realm sequence, the user is identified as Unknown.

You can optionally authenticate a user with the captive portal (that is, active authentication) if the user is not found in any realm in the sequence. For more information, see Captive Portal Guidelines and Limitations, on page 2526.

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Policies** > **Access Control heading** > **Identity** .

Step 3 Click **New Policy**.

Step 4 Enter a **Name** for the policy and an optional **Description**.

Step 5 Click **Save**.

Step 6 Click **Add Rule**.

Step 7 Enter a **Name** for the rule.

Step 8 From the list, click **Passive Authentication**.

Step 9 Click the **Realms & Settings** tab page.

Step 10 From the list, click the name of a realm or realm sequence.

The following figure shows an example.

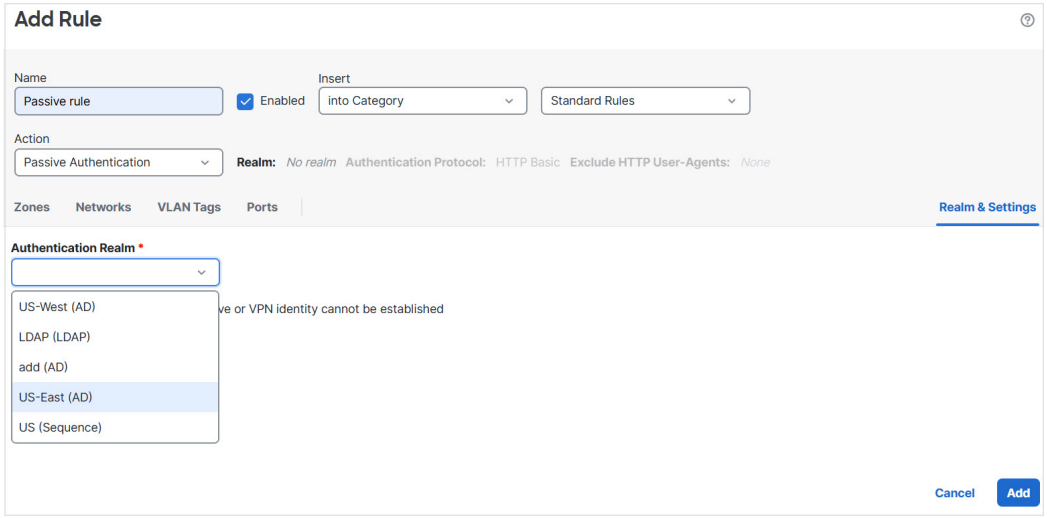

- If you choose a realm (such as **US-East** in the example), the system searches that realm for users to match the rule. If a user is not found, the user is identified as Unknown.
- If you choose a realm sequence (**US (Sequence)** in the example), users are searched in every realm in the sequence in the order specified in the sequence. If the user is not found, the user is identified as Unknown.
- You can also choose an LDAP realm.
- For additional ways to authenticate the user, check **Use active authentication if passive or VPN identity cannot be established**. For more information, see Captive Portal Guidelines and Limitations, on page 2526.

The following figure shows a sample passive identity policy configured to search for users in the US realm sequence.

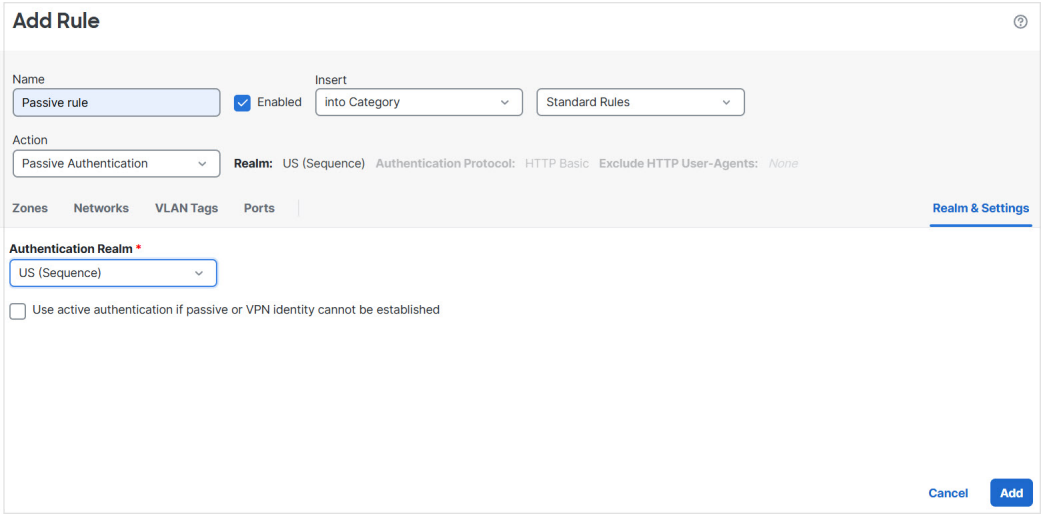

Step 11 (Optional.) To filter traffic by network object, click the **Identity Source** tab. From the list, click the network object to use to filter traffic for this identity policy. Click **Add** (+) to create a new network object.

Step 12 Set identity rule conditions as discussed in Identity Rule Conditions, on page 2557.

Step 13 Associate the identity rule with an access control rule as discussed in Associating Other Policies with Access Control, on page 1762.

Step 14 Deploy configuration changes to managed devices; see Deploy Configuration Changes, on page 259.

Create a Sample Identity Policy with an Active Authentication Rule

The associated tasks show an examples of configuring an identity policy with an **Active Authentication** rule where authentication is performed using either a realm or a realm sequence.

The differences follow:

- *Realm* enables you to use any supported authentication type (currently, **HTTP Basic**, **NTLM**, **Kerberos**, **HTTP Negotiate**, or **HTTP Response Page**).
- *Realm sequence* limits you to only the **HTTP Basic** or **HTTP Response Page** authentication types.

  Users who authenticate with a realm sequence and the HTTP Response Page authentication type see the following by default:

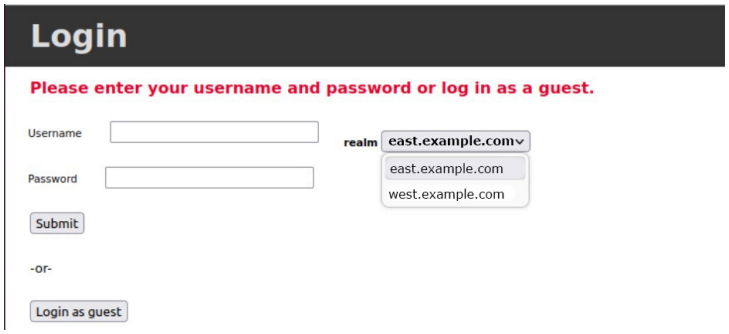

The user can authenticate in any of the following ways:

- If a list of realms in the realm sequence is displayed (as shown), the user must enter their user name and password in the provided fields and click the name of their realm from the list.
- If realms are not displayed in a list, the user can enter their credentials in `username@domain` format.

Users who authenticate with a realm and the HTTP Basic authentication page see the following:

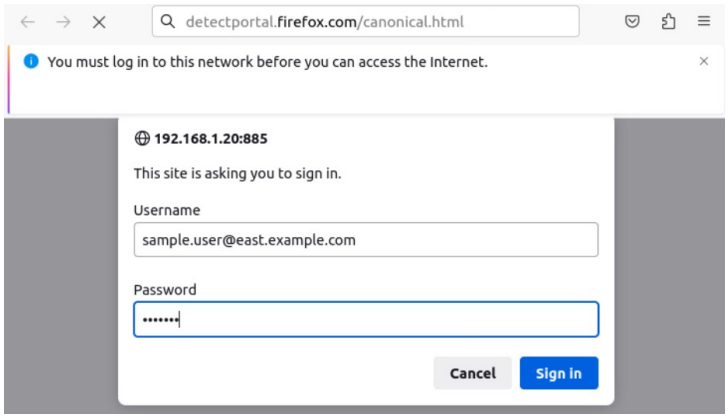

The user must enter their user name in `username@domain` format.

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Policies** > **Access Control heading** > **Identity** .

Step 3 Click **New Policy**.

Step 4 Enter a **Name** for the policy and an optional **Description**.

Step 5 Click **Save**.

Step 6 Click the **Active Authentication** tab.

Step 7 Enter the following information:

- **Server Certificate**: From the list, click an internal certificate object to use for secure connection to the Firewall Threat Defense device, or click **Add** (+) to add one.
- **Redirect to Host Name**: (Optional.) From the list, click a network object to which to redirect captive portal requests. If you omit this value, requests redirect to the IP address of the managed device. You can click **Add** (+) to create a new network object. For more information, see Redirect to Host Name Network Rule Conditions, on page 2558.

  Managed devices must have Snort 3 enabled to use this option.
- **Port**: Enter a port for the captive portal to use. This port must be unique for the captive portal and must match an access control rule you set up as discussed in Configure the Captive Portal Part 3: Create a TCP Port Access Control Rule, on page 2534. (Default is 885.)

- **Maximum login attempts**: Enter the maximum number of login attempts before login fails. (Default is 3.)
- **Share active authentication sessions across firewalls**: Uncheck to enable the Firewall Management Center to force users to reauthenticate every time they access your network using a different managed device than the last time.

  For more information about this option, see Captive Portal Fields, on page 2539.
- **Active Authentication Response Page**: Choose a system-provided or custom login page for captive portal users. For more information about your options, see Captive Portal Fields, on page 2539.

Step 8 Click **Save** to save your changes to the identity policy.

Step 9 Click the **Rules** tab.

Step 10 Click **Add Rule**.

Step 11 Enter a **Name** for the rule.

Step 12 From the list, click **Active Authentication**.

Step 13 Click the **Realms & Settings** tab page and continue with one of the following sections.

What to do next

Continue with one of the following sections:

- Active Authentication Using a Realm, on page 2570
- Active Authentication Using a Realm Sequence, on page 2571

Active Authentication Using a Realm

This task discusses how to authenticate captive portal users with a realm and any available authentication protocol (currently, **HTTP Basic**, **NTLM**, **Kerberos**, **HTTP Negotiate**, or **HTTP Response Page**).

Before you begin

Complete the task discussed in Create a Sample Identity Policy with an Active Authentication Rule, on page 2568.

Procedure

Step 1 Continue from Create a Sample Identity Policy with an Active Authentication Rule, on page 2568.

Step 2 On the **Realms & Settings** tab page, click **US-East**.

Step 3 From the **Authentication Protocol** list, click **NTLM**.

The following figure shows an example.

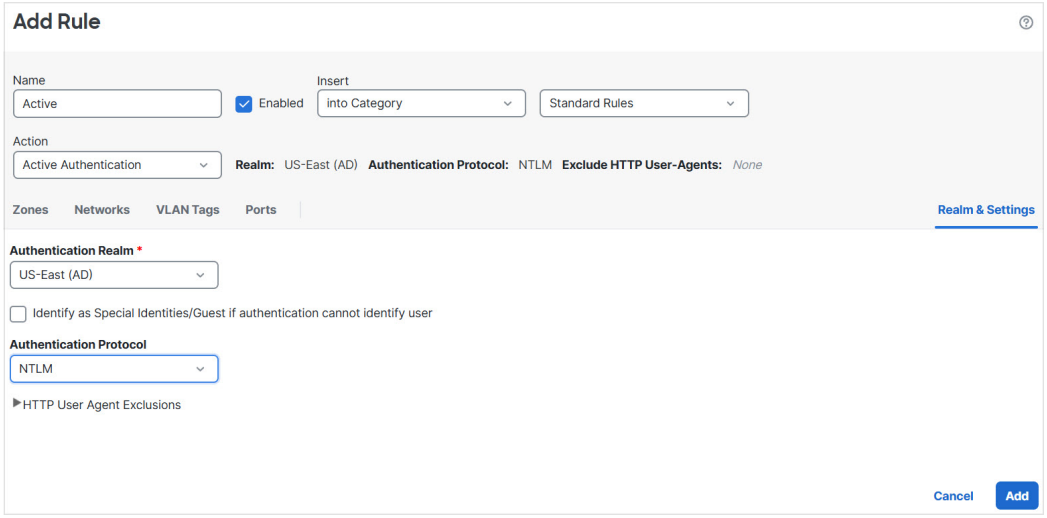

If you choose a realm (such as in the example), the system searches that realm for users to match the rule. If a user is not found, the user is identified as Unknown.

Step 4 Click **Add**.

Step 5 (Optional.) To filter traffic by network object, click the **Identity Source** tab. From the list, click the network object to use to filter traffic for this identity policy. Click **Add** (+) to create a new network object.

Step 6 Set identity rule conditions as discussed in Identity Rule Conditions, on page 2557.

Step 7 Associate the identity rule with an access control rule as discussed in Associating Other Policies with Access Control, on page 1762.

Step 8 Deploy configuration changes to managed devices; see Deploy Configuration Changes, on page 259.

Active Authentication Using a Realm Sequence

This task discusses how to authenticate captive portal users with a realm sequence, which restricts you to the **HTTP Basic** or **HTTP Response Page** authentication protocols.

Before you begin

Complete the task discussed in Create a Sample Identity Policy with an Active Authentication Rule, on page 2568.

Procedure

Step 1 Continue from Create a Sample Identity Policy with an Active Authentication Rule, on page 2568.

Step 2 On the **Realms & Settings** tab page, click the name of a realm from the list.

Step 3 From the list, click **US-East**.

Step 4 From the **Protocol** list, click **HTTP Response Page**.

The following figure shows an example.

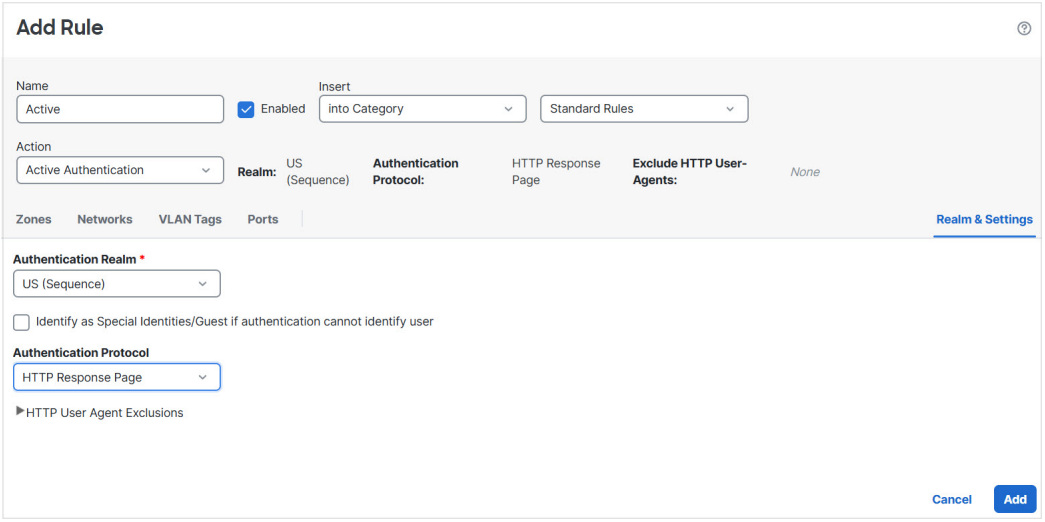

If you choose a realm sequence (such as in the example), the system searches realms in the sequence in the order specified in the realm sequence. The first realm in the sequence can be referred to as the *default* realm; it's the realm that is used if the user does not change it. If a user is not found, the user is identified as Unknown.

(Only if you upgraded to version 7.4.1 from an earlier version.) Edit the HTTP response page so it shows the list of realms in the sequence as discussed in Update a Custom Authentication Form, on page 2534.

Step 5 Click **Add**.

Step 6 (Optional.) To filter traffic by network object, click the **Identity Source** tab. From the list, click the network object to use to filter traffic for this identity policy. Click **Add** (+) to create a new network object.

Step 7 Set identity rule conditions as discussed in Identity Rule Conditions, on page 2557.

Step 8 Associate the identity rule with an access control rule as discussed in Associating Other Policies with Access Control, on page 1762.

Step 9 Deploy configuration changes to managed devices; see Deploy Configuration Changes, on page 259.

Manage an Identity Policy

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Policies** > **Access Control heading** > **Identity** .

Step 3 To delete a policy, click **Delete** (). If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 To edit a policy, click **Edit** () next to the policy and make changes as described in Create an Identity Policy, on page 2555. If **View** () appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 5 To copy a policy, click **Copy** ().

Step 6 To generate a report for the policy, click **Report** ( ) as described in Generate Current Policy Reports, on page 278.

Step 7 To compare policies, see Compare Policies, on page 277.

Step 8 To create a folder in which to organize policies, click **Add Category**.

What to do next

Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Manage an Identity Rule

Procedure

Step 1 Log in to the Firewall Management Center.

Step 2 Click **Policies** > **Access Control heading** > **Identity** .

Step 3 Click **Edit** ( ) next to the policy you want to edit. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 To edit an identity rule, click **Edit** ( ) and make changes as described in Create an Identity Policy, on page 2555.

Step 5 To delete an identity rule, click **Delete** ( ).

Step 6 To create a rule category, click **Add Category** and choose the position and the rule.

Step 7 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Troubleshoot User Control

If you notice unexpected user rule behavior, consider tuning your rule, identity source, or realm configurations. For other related troubleshooting information, see:

- Troubleshoot the ISE/ISE-PIC or Cisco TrustSec Issues, on page 2521
- Troubleshoot the TS Agent Identity Source, on page 2550
- Troubleshoot the Captive Portal Identity Source, on page 2541
- Troubleshoot Realms and User Downloads, on page 2456

Rules targeting realms, users, or user groups are not matching traffic

If you configure a TS Agent or ISE/ISE-PIC device to monitor a large number of user groups, or if you have a very large number of users mapped to hosts on your network, the system may drop user records due to your Firewall Management Center user limit. As a result, rules with user conditions may not match traffic as expected.

Rules targeting user groups or users within user groups are not matching traffic as expected

If you configure a rule with a user group condition, your LDAP or Active Directory server must have user groups configured. The system cannot perform user group control if the server organizes the users in basic object hierarchy.

Rules targeting users in secondary groups are not matching traffic as expected

If you configure a rule with a user group condition that includes or excludes users who are members of a secondary group on your Active Directory server, your server may be limiting the number of users it reports.

By default, Active Directory servers limit the number of users they report from secondary groups. You must customize this limit so that all of the users in your secondary groups are reported to the Firewall Management Center and eligible for use in rules with user conditions.

Rules are not matching users when seen for the first time

After the system detects activity from a previously-unseen user, the system retrieves information about them from the server. Until the system successfully retrieves this information, activity seen by this user is *not* handled by matching rules. Instead, the user session is handled by the next rule it matches (or the policy's default action, if applicable).

For example, this might explain when:

- Users who are members of user groups are not matching rules with user group conditions.
- Users who were reported by a TS Agent or ISE device are not matching rules, when the server used for user data retrieval is an Active Directory server.

Note that this might also cause the system to delay the display of user data in event views and analysis tools.

Rules are not matching all ISE/ISE-PIC users

This is expected behavior. You can perform user control on ISE/ISE-PIC users who were authenticated by an Active Directory domain controller. You cannot perform user control on ISE/ISE-PIC users who were authenticated by an LDAP, RADIUS, or RSA domain controller.

Users and groups using too much memory

If processing users and groups is using too much memory, health alerts are displayed. Remember that all user sessions are propagated to all devices managed by the Firewall Management Center. If your Firewall Management Center manages devices with different amounts of memory, the device with the least amount of memory determines the number of user sessions the system can handle without errors.

It's not possible to tune memory allocated to identity processes; even if a device has available memory, it can report out-of-memory issues. If issues persist, you have the following options:

- Segregate lower capacity managed devices on subnets and configure ISE/ISE-PIC to not report passive authentication data to those subnets.

See the chapter on managing network devices in the *Cisco Identity Services Engine Administrator Guide*.

- Unsubscribe from Security Group Tags (SGTs).

  For more information, see Ways to Configure the Cisco Identity Services Engine (Cisco ISE) Identity Source, on page 2512.

- Upgrade your managed device to a model with more memory.

PART XIII

Network Discovery

- Network Discovery Overview, on page 2579
- Host Identity Sources, on page 2587
- Application Detection, on page 2623
- Network Discovery Policies, on page 2643

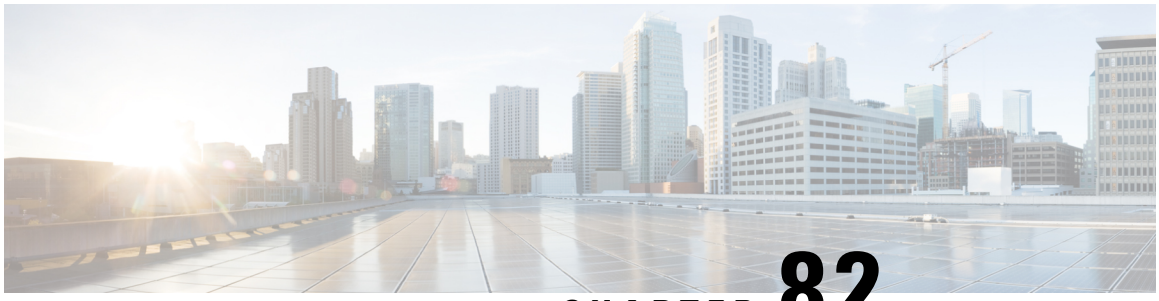

CHAPTER 82

Network Discovery Overview

The following topics discuss network discovery:

- About Detection of Host, Application, and User Data, on page 2579
- Host and Application Detection Fundamentals, on page 2580

About Detection of Host, Application, and User Data

The system uses *network discovery* and *identity* policies to collect host, application, and user data for traffic on your network. You can use certain types of discovery and identity data to build a comprehensive map of your network assets, perform forensic analysis, behavioral profiling, access control, and mitigate and respond to the vulnerabilities and exploits to which your organization is susceptible.

Host and Application Data

Host and application data is collected by host identity sources and application detectors according to the settings in your network discovery policy. Managed devices observe traffic on the network segments you specify.

For more information, see Host and Application Detection Fundamentals, on page 2580.

User Data

User data is collected by user identity sources according to the settings in your network discovery and identity policies. You can use the data for user awareness and user control.

For more information, see About User Identity, on page 2371.

Logging discovery and identity data allows you to take advantage of many features in the system, including:

- Viewing the network map, which is a detailed representation of your network assets and topology that you can view by grouping hosts and network devices, host attributes, application protocols, or vulnerabilities.
- Performing application and user control; that is, writing access control rules using application, realm, user, user group, and ISE attribute conditions.
- Viewing host profiles, which are complete views of all the information available for your detected hosts.
- Viewing dashboards, which (among other capabilities) can provide you with an at-a-glance view of your network assets and user activity.
- Viewing detailed information on the discovery events and user activity logged by the system.

- Associating hosts and any servers or clients they are running with the exploits to which they are susceptible.

  This enables you to identify and mitigate vulnerabilities, evaluate the impact that intrusion events have on your network, and tune intrusion rule states so that they provide maximum protection for your network assets

- Alerting you by email, SNMP trap, or syslog when the system generates either an intrusion event with a specific impact flag, or a specific type of discovery event
- Monitoring your organization's compliance with an allow list of allowed operating systems, clients, application protocols, and protocols
- Creating correlation policies with rules that trigger and generate correlation events when the system generates discovery events or detects user activity
- Logging and using NetFlow connections, if applicable.

Host and Application Detection Fundamentals

You can configure your network discovery policy to perform host and application detection.

For more information, see Overview: Host Data Collection, on page 2587 and Overview: Application Detection, on page 2623.

Passive Detection of Operating System and Host Data

Passive detection is the system's default method of populating the network map by analyzing network traffic (and any exported NetFlow data). Passive detection provides contextual information about your network assets, such as operating systems and running applications.

If traffic from a monitored host does not offer conclusive evidence of the host's operating system, the network map displays the most likely operating system. For example, a NAT device may appear to be running several operating systems because of the hosts "behind" the NAT device. To make this most-likely determination, the system uses a confidence value it assigns to each detected operating system, and the amount of corroborating data among detected operating systems.

Note The system does not consider reported "`unknown`" applications and operating systems in its determination.

If passive detection inaccurately identifies your network assets, consider the placement of your managed devices. You can also augment the system's passive detection capabilities with custom operating-system fingerprints and custom application detectors. Or, you can use *active detection*, which is not based on traffic analysis, but instead allows you to directly update the network map using scan results or other information sources.

Active Detection of Operating System and Host Data

Active detection adds host information collected by active sources to network maps. For example, you can use the Nmap scanner to actively scan the hosts that you target on your network. Nmap discovers operating systems and applications on hosts.

In addition, the host input feature allows you to actively add *host input data* to network maps. There are two different categories of host input data:

- *user input data*—Data added through the system user interface. You can modify a host's operating system or application identity through this interface.
- *host import input data*—Data imported using a command line utility.

The system retains one identity for each active source. When you run an Nmap scan instance, for example, the results of the previous scan are replaced with the new scan results. However, if you run an Nmap scan and then replace those results with data from a client whose results are imported through the command line, the system retains both the identities from the Nmap results and the identities from the import client. The system then uses the priorities set in the network discovery policy to determine which active identity to use as the current identity.

Note that user input is considered one source, even if it comes from different users. As an example, if UserA sets the operating system through the host profile, and then UserB changes that definition through the host profile, the definition set by UserB is retained, and the definition set by UserA is discarded. In addition, note that user input overrides all other active sources and is used as the current identity if it exists.

Current Identities for Applications and Operating Systems

The *current identity* for an application or an operating system on a host is the identity that the system finds most likely to be correct.

The system uses the current identity for an operating system or application for the following purposes:

- to assign vulnerabilities to a host
- for impact assessment
- when evaluating correlation rules written against operating system identifications, host profile qualifications, and compliance allow lists
- for display in the Hosts and Servers table views in workflows
- for display in the host profile
- to calculate the operating system and application statistics on the Discovery Statistics page

The system uses source priorities to determine which active identity should be used as the current identity for an application or operating system.

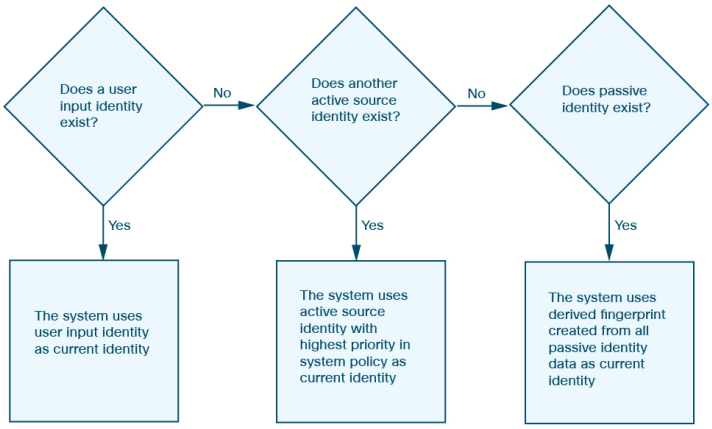

For example, if a user sets the operating system to Windows 2003 Server on a host, Windows 2003 Server is the current identity. Attacks which target Windows 2003 Server vulnerabilities on that host are given a higher impact, and the vulnerabilities listed for that host in the host profile include Windows 2003 Server vulnerabilities.

The database may retain information from several sources for the operating system or for a particular application on a host.

The system treats an operating system or application identity as the current identity when the source for the data has the highest source priority. Possible sources have the following priority order:

1. user

2. scanner and application (set in the network discovery policy)

3. managed devices

4. NetFlow records

A new higher priority application identity will not override a current application identity if it has less detail than the current identity.

In addition, when an identity conflict occurs, the resolution of the conflict depends on settings in the network discovery policy or on your manual resolution.

Current User Identities

When the system detects multiple logins to the same host by different users, the system assumes that only one user is logged into any given host at a time, and that the current user of a host is the last authoritative user login. If only non-authoritative user logins have been logged into the host, the last non-authoritative user login is considered the current user. If multiple users are logged in through remote sessions, the last user reported by the server is the user reported to the Firewall Management Center.

When the system detects multiple logins to the same host by the same user, the system records the first time that a user logs into a specific host and disregards subsequent logins. If an individual user is the only person who logs into a specific host, the only login that the system records is the original login.

If another user logs into that host, however, the system records the new login. Then, if the original user logs in again, his or her new login is recorded.

Application and Operating System Identity Conflicts

An *identity conflict* occurs when the system reports a new passive identity that conflicts with the current active identity and previously reported passive identities. For example, the previous passive identity for an operating system is reported as Windows 2000, then an active identity of Windows XP becomes current. Next, the system detects a new passive identity of Ubuntu Linux 8.04.1. The Windows XP and the Ubuntu Linux identities are in conflict.

When an identity conflict exists for the identity of the host's operating system or one of the applications on the host, the system lists both conflicting identities as current and uses both for impact assessment until the conflict is resolved.

A user with Administrator privileges can resolve identity conflicts automatically by choosing to always use the passive identity or always use the active identity. Unless you disable automatic resolution of identity conflicts, identity conflicts are always automatically resolved.

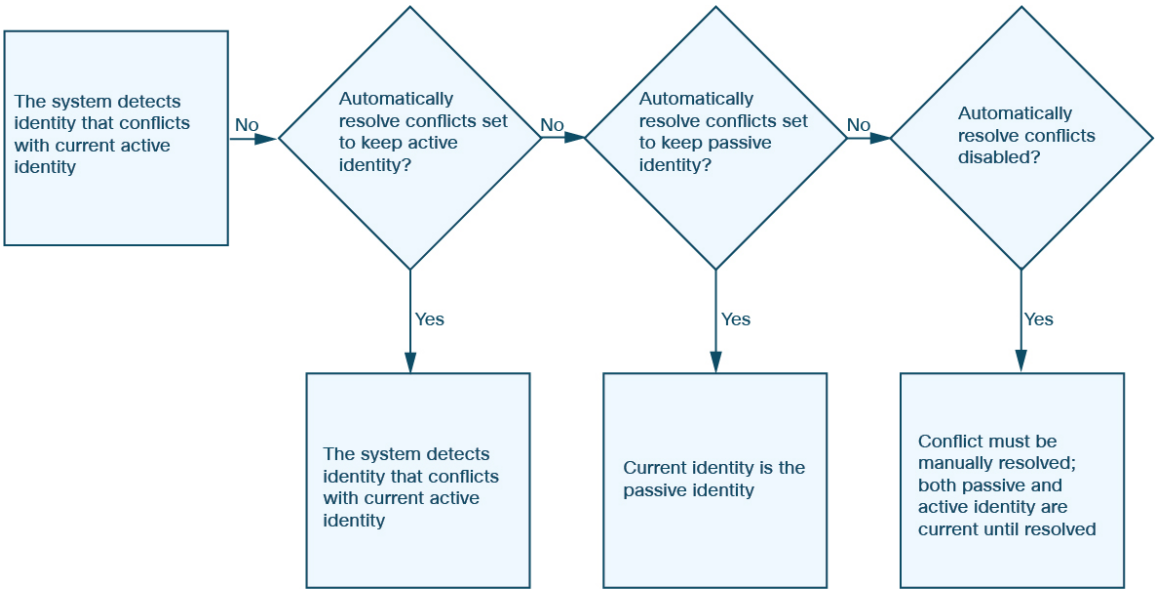

A user with Administrator privileges can also configure the system to generate an event when an identity conflict occurs. That user can then set up a correlation policy with a correlation rule that uses an Nmap scan as a correlation response. When an event occurs, Nmap scans the host to obtain updated host operating system and application data.

NetFlow Data

NetFlow is a Cisco IOS application that provides statistics on packets flowing through a router. It is available on Cisco networking devices and can also be embedded in Juniper, FreeBSD, and OpenBSD devices.

When NetFlow is enabled on a network device, a database on the device (the NetFlow cache) stores records of the flows that pass through the router. A flow, called a *connection* in the system, is a sequence of packets that represents a session between a source and destination host, using specific ports, protocol, and application protocol. The network device can be configured to export this NetFlow data. In this documentation, network devices configured in this way are called *NetFlow exporters*.

Managed devices can be configured to collect records from NetFlow exporters, generate unidirectional end-of-connection events based on the data in those records, and finally send those events to the Firewall

Management Center to be logged in the connection event database. You can also configure the network discovery policy to add host and application protocol information to the database based on the information in NetFlow connections.

You can use this discovery and connection data to supplement the data gathered directly by your managed devices. This is especially useful if you have NetFlow exporters monitoring networks that your managed devices cannot monitor.

Requirements for Using NetFlow Data

Before you configure the system to analyze NetFlow data, you must enable the NetFlow feature on the routers or other NetFlow-enabled network devices you plan to use, and configure the devices to broadcast NetFlow data to a destination network where the sensing interface of a managed device is connected.

The system can parse both NetFlow version 5 and NetFlow version 9 records. NetFlow exporters **must** use one of those versions if you want to export the data to the system. In addition, the system requires that specific fields be present in the exported NetFlow templates and records. If your NetFlow exporters are using version 9, which you can customize, you **must** make sure that the exported templates and records contain the following fields, in any order:

- IN_BYTES (1)
- IN_PKTS (2)
- PROTOCOL (4)
- TCP_FLAGS (6)
- L4_SRC_PORT (7)
- IPV4_SRC_ADDR (8)
- L4_DST_PORT (11)
- IPV4_DST_ADDR (12)
- LAST_SWITCHED (21)
- FIRST_SWITCHED (22)
- IPV6_SRC_ADDR (27)
- IPV6_DST_ADDR (28)

Because the system uses managed devices to analyze NetFlow data, your deployment must include at least one managed device that can monitor NetFlow exporters. At least one sensing interface on that managed device must be connected to a network where it can collect the exported NetFlow data. Because the sensing interfaces on managed devices do not usually have IP addresses, the system does not support the direct collection of NetFlow records.

Note that the Sampled NetFlow feature available on some network devices collects NetFlow statistics on only a subset of packets that pass through the devices. Although enabling this feature can improve CPU utilization on the network device, it may affect the NetFlow data you are collecting for analysis by the system.

Differences between NetFlow and Managed Device Data

The traffic represented by NetFlow data is not directly analyzed. Instead, it converts exported NetFlow records into connection logs and host and application protocol data.

As a result, there are several differences between converted NetFlow data and the discovery and connection data gathered directly by your managed devices. You should keep these differences in mind when performing analysis that requires:

- Statistics on the number of detected connections
- Operating system and other host-related information (including vulnerabilities)
- Application data, including client information, web application information, and vendor and version server information
- Knowing which host in a connection is the initiator and which is the responder

Network Discovery Policy versus Access Control Policy

You configure NetFlow data collection, including connection logging, using rules in the network discovery policy. Contrast this with connection logging for connections detected by managed devices, which you configure per access control rule.

Types of Connection Events

Because NetFlow data collection is linked to networks rather than access control rules, you do not have granular control over which NetFlow connections the system logs.

NetFlow data cannot generate Security Intelligence events.

NetFlow-based connection events can be stored in the connection event database only; you cannot send them to the system log or an SNMP trap server.

Number of Connection Events Generated Per Monitored Session

For connections detected directly by managed devices, you can configure the access control rule to log a bidirectional connection event at the beginning or end of a connection, or both.

In contrast, because exported NetFlow records contain unidirectional connection data, the system generates at least two connection events for each NetFlow record it processes. This also means that a summary's connection count is incremented by two for every connection based on NetFlow data, providing an inflated count of the number of connections that are actually occurring on your network.

Because the NetFlow exporter outputs records at a fixed interval even if a connection is still ongoing, long-running sessions can result in multiple exported records, each of which generates a connection event. For example, if the NetFlow exporter exports every five minutes, and a particular connection lasts twelve minutes, the system generates six connection events for that session:

- One pair of events for the first five minutes
- One pair for the second five minutes
- A final pair when the connection is terminated

Host and Operating System Data

Hosts added to the network map from NetFlow data do not have operating system, NetBIOS, or host type (host vs network device) information. You can, however, manually set a host’s operating system identity using the host input feature.

Application Data

For connections detected directly by managed devices, the system can identify application protocols, clients, and web applications by examining the packets in the connection.

When the system processes NetFlow records, the system uses a port correlation in `/etc/sf/services` to extrapolate application protocol identity. However, there is no vendor or version information for those application protocols, nor do connection logs contain information on client or web applications used in the session. You can, however, manually provide this information using the host input feature.

Note that a simple port correlation means that application protocols running on non-standard ports may be unidentified or misidentified. Additionally, if no correlation exists, the system marks the application protocol as `unknown` in connection logs.

Vulnerability Mappings

The system cannot map vulnerabilities to hosts monitored by NetFlow exporters, unless you use the host input feature to manually set either a host's operating system identity or an application protocol identity. Note that because there is no client information in NetFlow connections, you cannot associate client vulnerabilities with hosts created from NetFlow data.

Initiator and Responder Information in Connections

For connections detected directly by managed devices, the system can identify which host is the initiator, or source, and which is the responder, or destination. However, NetFlow data does not contain initiator or responder information.

When the System processes NetFlow records, it uses an algorithm to determine this information based on the ports each host is using, and whether those ports are well-known:

- If both or neither port being used is a well-known port, the system considers the host using the lower-number port to be the responder.
- If only one of the hosts is using a well-known port, the system considers that host to be the responder.

For this purpose, a well-known port is any port that is either numbered from 1 to 1023, or that contains application protocol information in `/etc/sf/services` on the managed device.

In addition, for connections detected directly by managed devices, the system records two byte counts in the corresponding connection event:

- The **Initiator Bytes** field records bytes sent.
- The **Responder Bytes** field records bytes received.

Connection events based on unidirectional NetFlow records contain only one byte count, which the system assigns to either **Initiator Bytes** or **Responder Bytes**, depending on the port-based algorithm. The system sets the other field to 0. Note that if you are viewing connection summaries (aggregated connection data) of NetFlow records, both fields may be populated.

NetFlow-only Connection Event Fields

A small number of fields are present only in connection events generated from NetFlow records; see *Information Available in Connection Event Fields* in the Cisco Secure Firewall Management Center Administration Guide.

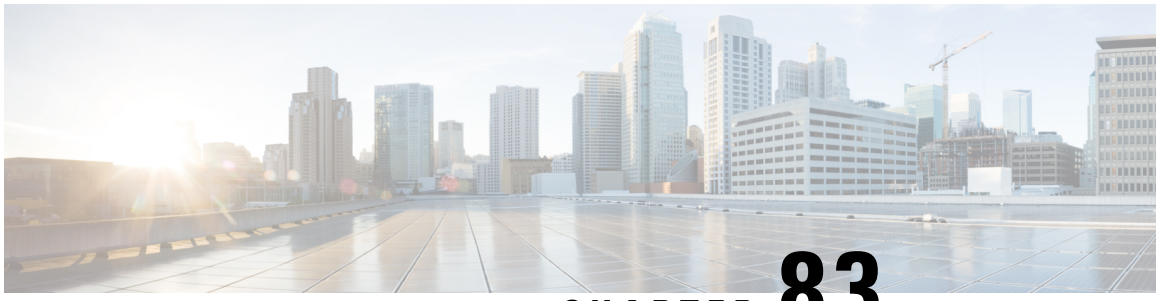

CHAPTER 83

Host Identity Sources

The following topics provide information on host identity sources:

- Overview: Host Data Collection, on page 2587
- Requirements and Prerequisites for Host Identity Sources, on page 2588
- Determining Which Host Operating Systems the System Can Detect, on page 2588
- Identifying Host Operating Systems, on page 2588
- Custom Fingerprinting, on page 2589
- Host Input Data, on page 2597
- Nmap Scanning, on page 2604
- History for Host Identity Sources, on page 2622

Overview: Host Data Collection

As the system passively monitors the traffic that travels through your network, it compares specific packet header values and other unique data from network traffic against established definitions (called *fingerprints*) to determine information about the hosts on your network, including:

- the number and types of hosts (including network devices such as bridges, routers, load balancers, and NAT devices)
- basic network topology data, including the number of hops from the discovery point on the network to the hosts
- the operating systems running on the hosts
- applications on the hosts and users associated with these applications

If the system cannot identify a host's operating system, you can create custom client or server fingerprints. The system uses these fingerprints to identify new hosts. You can map fingerprints to systems in the vulnerability database (VDB) to allow the appropriate vulnerability information to be displayed whenever a host is identified using the custom fingerprint.

Note In addition to collecting host data from monitored network traffic, the system can collect host data from exported NetFlow records, and you can actively add host data using Nmap scans and the host input feature.

Requirements and Prerequisites for Host Identity Sources

Model Support

Any.

Supported Domains

Any, with the exception of custom fingerprinting, which is Leaf only.

User Roles

- Admin
- Discovery Admin, except for third-party data and custom mappings.

Determining Which Host Operating Systems the System Can Detect

To learn which exact operating systems the system can fingerprint, view the list of available fingerprints that is shown during the process of creating a custom OS fingerprint.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Custom Operating Systems**.

Step 3 Click **Create Custom Fingerprint**.

Step 4 View the lists of options in the drop-down lists in the **OS Vulnerability Mappings** section. These options are the operating systems that the system can fingerprint.

What to do next

As needed, see Identifying Host Operating Systems, on page 2588.

Identifying Host Operating Systems

If the system does not correctly identify a host's operating system (for example, it shows in the Host Profile as Unknown or is incorrectly identified), try the strategies below.

Procedure

Try one of the following strategies:

- Check the Network Discovery Identity Conflict Settings.
- Create a custom fingerprint for the host.
- Run an Nmap scan against the host.
- Import data into the network map, using the host input feature.
- Manually enter operating system information.

Custom Fingerprinting

The system includes operating system *fingerprints* that the system uses to identify the operating system on each host it detects. However, sometimes the system cannot identify a host operating system or misidentifies it because no fingerprints exist that match the operating system. To correct this problem, you can create a *custom fingerprint,* which provides a pattern of operating system characteristics unique to the unknown or misidentified operating system, to supply the name of the operating system for identification purposes.

If the system cannot match a host's operating system, it cannot identify the vulnerabilities for the host, because the system derives the list of vulnerabilities for each host from its operating system fingerprint. For example, if the system detects a host running Microsoft Windows, the system has a stored Microsoft Windows vulnerability list that it adds to the host profile for that host based on the detected Windows operating system.

As an example, if you have several devices on your network running a new beta version of Microsoft Windows, the system cannot identify that operating system or map vulnerabilities to the hosts. However, knowing that the system has a list of vulnerabilities for Microsoft Windows, you may want to create a custom fingerprint for one of the hosts to help identify the other hosts running the same operating system. You can include a mapping of the vulnerability list for Microsoft Windows in the fingerprint to associate that list with each host that matches the fingerprint.

When you create a custom fingerprint, the Firewall Management Center lists the set of vulnerabilities associated with that fingerprint for any hosts running the same operating system. If the custom fingerprint you create does not have any vulnerabilities mappings in it, the system uses the fingerprint to assign the custom operating system information you provide in the fingerprint. When the system sees new traffic from a previously detected host, the system updates the host with the new fingerprint information. the system also uses the new fingerprint to identify any new hosts with that operating system the first time they are detected.

Before creating a custom fingerprint, you should determine why the host is not being identified correctly to decide whether custom fingerprinting is a viable solution.

You can create two types of fingerprints with the system:

- Client fingerprints, which identify operating systems based on the SYN packet that the host sends when it connects to a TCP application running on another host on the network.

- Server fingerprints, which identify operating systems based on the SYN-ACK packet that the host uses to respond to an incoming connection to a running TCP application.

Note If both a client and server fingerprint match the same host, the client fingerprint is used.

After creating fingerprints, you must activate them before the system can associate them with hosts.

Related Topics

Creating a Custom Fingerprint for Clients, on page 2592
Creating a Custom Fingerprint for Servers, on page 2594

Managing Fingerprints

After a fingerprint is created and activated, you can edit a fingerprint to make changes or add vulnerability mappings.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Custom Operating Systems**. If the system is awaiting data to create a fingerprint, it automatically refreshes the page every 10 seconds until the fingerprint is created.

Step 3 Manage your custom fingerprints:

- Activate/Deactivate — Activate or deactivate a fingerprint as described in Activating and Deactivating Fingerprints, on page 2590.
- Create — Create fingerprints as described in Creating a Custom Fingerprint for Clients, on page 2592 and Creating a Custom Fingerprint for Servers, on page 2594.
- Edit — Edit a fingerprint as described in Editing an Active Fingerprint, on page 2591 and Editing an Inactive Fingerprint, on page 2591.
- Delete — Click **Delete** ( ) next to the fingerprint you want to delete, and click **OK** to confirm. You can only delete deactivated fingerprints.

Activating and Deactivating Fingerprints

You must activate a custom fingerprint before the system can use it to identify hosts. After the new fingerprint is activated, the system uses it to re-identify previously discovered hosts and discover new hosts.

If you want to stop using a fingerprint, you can deactivate it. Deactivating a fingerprint causes a fingerprint to no longer be used, but allows it to remain on the system. When you deactivate a fingerprint, the operating system is marked as unknown for hosts that use the fingerprint. If the hosts are detected again and match a different active fingerprint, they are then identified by that active fingerprint.

Deleting a fingerprint removes it from the system completely. After deactivating a fingerprint, you can delete it.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Custom Operating Systems**.

Step 3 Click the slider next to the fingerprint you want to activate or deactivate.

Note
The activate option is only available if the fingerprint you created is valid. If the slider is not available, try creating the fingerprint again.

Editing an Active Fingerprint

If a fingerprint is *active*, you can modify the fingerprint name, description, custom operating system display, and map additional vulnerabilities to it.

You can modify the fingerprint name, description, custom operating system display, and map additional vulnerabilities to it.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Custom Operating Systems**

Step 3 Click **Edit** (✎) next to the fingerprint you want to edit.

Step 4 Modify the fingerprint name, description, and custom OS display, if necessary.

Step 5 If you want to delete a vulnerability mapping, click **Delete** next to the mapping in the **Pre-Defined OS Product Maps** section of the page.

Step 6 If you want to add additional operating systems for vulnerability mapping, choose the **Product** and, if applicable, **Major Version**, **Minor Version**, **Revision Version**, **Build**, **Patch**, and **Extension** and then click **Add OS Definition**.

The vulnerability mapping is added to the **Pre-Defined OS Product Maps** list.

Step 7 Click **Save**.

Editing an Inactive Fingerprint

If a fingerprint is *inactive*, you can modify all elements of the fingerprint and resubmit it to the Secure Firewall Management Center. This includes all properties you specified when creating the fingerprint, such as fingerprint type, target IP addresses and ports, vulnerability mappings, and so on. When you edit an inactive fingerprint and submit it, it is resubmitted to the system and, if it is a client fingerprint, you must resend traffic to the appliance before activating it. Note that you can choose only a single vulnerability mapping for an inactive fingerprint. After you activate the fingerprint, you can map additional operating systems and versions to its vulnerabilities list.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Custom Operating Systems**.

Step 3 Click **Edit** (✎) next to the fingerprint you want to edit.

Step 4 Make changes to the fingerprint as necessary:

- If you are modifying a client fingerprint, see Creating a Custom Fingerprint for Clients, on page 2592.
- If you are modifying a server fingerprint, see Creating a Custom Fingerprint for Servers, on page 2594.

Step 5 Click **Save**.

What to do next

- If you modified a client fingerprint, remember to send traffic from the host to the appliance gathering the fingerprint.

Creating a Custom Fingerprint for Clients

Client fingerprints identify operating systems based on the SYN packet a host sends when it connects to a TCP application running on another host on the network.

If the Firewall Management Center does not have direct contact with monitored hosts, you can specify a device that is managed by the Firewall Management Center and is closest to the host you intend to fingerprint when specifying client fingerprint properties.

Before you begin the fingerprinting process, obtain the following information about the host you want to fingerprint:

- The number of network hops between the host and the Firewall Management Center or the device you use to obtain the fingerprint. (Cisco strongly recommends that you directly connect the Firewall Management Center or the device to the same subnet that the host is connected to.)
- The network interface (on the Firewall Management Center or the device) that is connected to the network where the host resides.
- The actual operating system vendor, product, and version of the host.
- Access to the host in order to generate client traffic.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Custom Operating Systems**.

Step 3 Click **Create Custom Fingerprint**.

Step 4 From the **Device** drop-down list, choose the Firewall Management Center or the device that you want to use to collect the fingerprint.

Step 5 Enter a **Fingerprint Name**.

Step 6 Enter a **Fingerprint Description**.

Step 7 From the **Fingerprint Type** list, choose **Client**.

Step 8 In the **Target IP Address** field, enter an IP address of the host you want to fingerprint.

Note that the fingerprint will only be based on traffic to and from the host IP address you specify, not any of the host's other IP addresses (if it has any).

Step 9 In the **Target Distance** field, enter the number of network hops between the host and the device that you chose earlier to collect the fingerprint.

Caution
This must be the actual number of physical network hops to the host, which may or may not be the same as the number of hops detected by the system.

Step 10 From the **Interface** list, choose the network interface that is connected to the network segment where the host resides.

Note
Use the management interface or any other available network interfaces that is configured as the sensing interface on your device to perform fingerprint collection. If you do not know which interface is the sensing interface on your device, refer to the *Installation Guide* for the specific model you are using to fingerprint.

Step 11 If you want to display custom information in the host profile for fingerprinted hosts (or if the host you want to fingerprint does not reside in the **OS Vulnerability Mappings** section), choose **Use Custom OS Display** and provide the values you want to display for the following:

- In the **Vendor String** field, enter the operating system's vendor name. For example, the vendor for Microsoft Windows would be Microsoft.
- In the **Product String** field, enter the operating system's product name. For example, the product name for Microsoft Windows 2000 would be Windows.
- In the **Version String** field, enter the operating system's version number. For example, the version number for Microsoft Windows 2000 would be 2000.

Step 12 In the OS Vulnerability Mappings section, choose the operating system, product, and versions you want to use for vulnerability mapping.

You must specify **Vendor** and **Product** values in this section if you want to use the fingerprint to identify vulnerabilities for matching hosts or if you do not assign custom operating system display information.

To map vulnerabilities for all versions of an operating system, specify only the **Vendor** and **Product** values.

Note
Not all options in the **Major Version**, **Minor Version**, **Revision Version**, **Build**, **Patch**, and **Extension** drop-down lists may apply to the operating system you choose. In addition, if no definition appears in a list that matches the operating system you want to fingerprint, you can leave these values empty. Be aware that if you do not create any OS vulnerability mappings in a fingerprint, the system cannot use the fingerprint to assign a vulnerabilities list with hosts identified by the fingerprint.

Example:

If you want your custom fingerprint to assign the list of vulnerabilities from Redhat Linux 9 to matching hosts, choose **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the major version.

Example:

To add all versions of the Palm OS, you would choose **PalmSource, Inc.** from the **Vendor** list, **Palm OS** from the **Product** list, and leave all other lists at their default settings.

Step 13 Click **Create**.

The status briefly shows `New`, then switches to `Pending`, where it remains until traffic is seen for the fingerprint. Once traffic is seen, it switches to `Ready`.

The Custom Fingerprint status page refreshes every ten seconds until it receives data from the host in question.

Step 14 Using the IP address you specified as the target IP address, access the host you are trying to fingerprint and initiate a TCP connection to the appliance.

To create an accurate fingerprint, traffic **must** be seen by the appliance collecting the fingerprint. If you are connected through a switch, traffic to a system other than the appliance may not be seen by the system.

Example:

Access the web interface of the Firewall Management Center from the host you want to fingerprint or SSH into the Firewall Management Center from the host. If you are using SSH, use the command below, where localIPv6address is the IPv6 address specified in step 7 that is currently assigned to the host and DCmanagementIPv6address is the management IPv6 address of the Firewall Management Center. The Custom Fingerprint page should then reload with a “Ready” status.

```
ssh -b localIPv6address DCmanagementIPv6address
```

What to do next

- Activate the fingerprint as described in Activating and Deactivating Fingerprints, on page 2590.

Creating a Custom Fingerprint for Servers

Server fingerprints identify operating systems based on the SYN-ACK packet that the host uses to respond to an incoming connection to a running TCP application. Before you begin, you should obtain the following information about the host you want to fingerprint:

- The number of network hops between the host and the appliance you use to obtain the fingerprint. Cisco strongly recommends that you directly connect an unused interface on the appliance to the same subnet that the host is connected to.
- The network interface (on the appliance) that is connected to the network where the host resides.
- The actual operating system vendor, product, and version of the host.
- An IP address that is not currently in use and is authorized on the network where the host is located.

Tip If the Firewall Management Center does not have direct contact with monitored hosts, you can specify a managed device that is closest to the host you intend to fingerprint when specifying server fingerprint properties.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Custom Operating Systems**.

Step 3 Click **Create Custom Fingerprint**.

Step 4 From the **Device** list, choose the Firewall Management Center or the managed device that you want to use to collect the fingerprint.

Step 5 Enter a **Fingerprint Name**.

Step 6 Enter a **Fingerprint Description**.

Step 7 From the **Fingerprint Type** list, choose **Server** to display the server fingerprinting options.

Step 8 In the **Target IP Address** field, enter an IP address of the host you want to fingerprint.

Note that the fingerprint will only be based on traffic to and from the host IP address you specify, not any of the host's other IP addresses (if it has any).

Caution
You can capture IPv6 fingerprints only with appliances running Version 5.2 and later.

Step 9 In the **Target Distance** field, enter the number of network hops between the host and the device that you chose earlier to collect the fingerprint.

Caution
This must be the actual number of physical network hops to the host, which may or may not be the same as the number of hops detected by the system.

Step 10 From the **Interface** list, choose the network interface that is connected to the network segment where the host resides.

Caution
Cisco recommends that you do **not** use the sensing interface on a managed device for fingerprinting for several reasons. First, fingerprinting does not work if the sensing interface is on a span port. Also, if you use the sensing interface on a device, the device stops monitoring the network for the amount of time it takes to collect the fingerprint. You can, however, use the management interface or any other available network interfaces to perform fingerprint collection. If you do not know which interface is the sensing interface on your device, refer to the *Installation Guide* for the specific model you are using to fingerprint.

Step 11 Click **Get Active Ports**.

Step 12 In the **Server Port** field, enter the port that you want the device chose to collect the fingerprint to initiate contact with, or choose a port from the **Get Active Ports** drop-down list.

You can use any server port that you know is open on the host (for instance, 80 if the host is running a web server).

Step 13 In the **Source IP Address** field, enter an IP address that should be used to attempt to communicate with the host.

You should use a source IP address that is authorized for use on the network but is not currently being used, for example, a DHCP pool address that is currently not in use. This prevents you from temporarily knocking another host offline while you create the fingerprint.

You should exclude that IP address from monitoring in your network discovery policy while you create the fingerprint. Otherwise, the network map and discovery event views will be cluttered with inaccurate information about the host represented by that IP address.

Step 14 In the **Source Subnet Mask** field, enter the subnet mask for the IP address you are using.

Step 15 If the **Source Gateway** field appears, enter the default gateway IP address that should be used to establish a route to the host.

Step 16 If you want to display custom information in the host profile for fingerprinted hosts or if the fingerprint name you want to use does not exist in the OS Definition section, choose **Use Custom OS Display** in the Custom OS Display section.

Provide the values you want to appear in host profiles for the following:

- In the **Vendor String** field, enter the operating system's vendor name. For example, the vendor for Microsoft Windows would be Microsoft.
- In the **Product String** field, enter the operating system's product name. For example, the product name for Microsoft Windows 2000 would be Windows.
- In the **Version String** field, enter the operating system's version number. For example, the version number for Microsoft Windows 2000 would be 2000.

Step 17 In the OS Vulnerability Mappings section, choose the operating system, product, and versions you want to use for vulnerability mapping.

You must specify a Vendor and Product name in this section if you want to use the fingerprint to identify vulnerabilities for matching hosts or if you do not assign custom operating system display information.

To map vulnerabilities for all versions of an operating system, specify only the vendor and product name.

Note
Not all options in the **Major Version**, **Minor Version**, **Revision Version**, **Build**, **Patch**, and **Extension** drop-down lists may apply to the operating system you choose. In addition, if no definition appears in a list that matches the operating system you want to fingerprint, you can leave these values empty. Be aware that if you do not create any OS vulnerability mappings in a fingerprint, the system cannot use the fingerprint to assign a vulnerabilities list with hosts identified by the fingerprint.

Example:

If you want your custom fingerprint to assign the list of vulnerabilities from Redhat Linux 9 to matching hosts, choose **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.

Example:

To add all versions of the Palm OS, you would choose **PalmSource, Inc.** from the **Vendor** list, **Palm OS** from the **Product** list, and leave all other lists at their default settings.

Step 18 Click **Create**.
The Custom Fingerprint status page refreshes every ten seconds and should reload with a "Ready" status.

Note
If the target system stops responding during the fingerprinting process, the status shows an `ERROR: No Response` message. If you see this message, submit the fingerprint again. Wait three to five minutes (the time period may vary depending on the target system), click **Edit** () to access the Custom Fingerprint page, and then click **Create**.

What to do next

- Activate the fingerprint as described in Activating and Deactivating Fingerprints, on page 2590.

Host Input Data

You can augment the network map by importing network map data from third parties. You can also use the host input feature by modifying operating system or application identities or deleting application protocols, protocols, host attributes, or clients using the web interface.

The system may reconcile data from multiple sources to determine the current identity of an operating system or application.

All data except third-party vulnerabilities is discarded when the affected host is removed from the network map. For more information on setting up scripts or import files, see the *Firepower System Host Input API Guide*.

To include imported data in impact correlations, you must map the data to the operating system and application definitions in the database.

Requirements for Using Third-Party Data

You can import discovery data from third-party systems on your network. However, to enable features where intrusion and discovery data are used together, such as Cisco recommendations, adaptive profile updates, or impact assessment, you should map as many elements of it as possible to corresponding definitions. Consider the following requirements for using third-party data:

- If you have a third-party system that has specific data on your network assets, you can import that data using the host input feature. However, because third parties may name the products differently, you must map the third-party vendor, product, and versions to the corresponding Cisco product definition. After you map the products, you must enable vulnerability mappings for impact assessment in the Firewall Management Center configuration to allow impact correlation. For versionless or vendorless application protocols, you need to map vulnerabilities for the application protocols in the Firewall Management Center configuration.
- If you import patch information from a third party and you want to mark all vulnerabilities fixed by that patch as invalid, you must map the third-party fix name to a fix definition in the database. All vulnerabilities addressed by the fix will then be removed from hosts where you add that fix.
- If you import operating system and application protocol vulnerabilities from a third party and you want to use them for impact correlation, you must map the third-party vulnerability identification string to vulnerabilities in the database. Note that although many clients have associated vulnerabilities, and clients are used for impact assessment, you cannot import and map third-party client vulnerabilities. After the vulnerabilities are mapped, you must enable third-party vulnerability mappings for impact assessment in the Firewall Management Center configuration. To cause application protocols without vendor or version information to map to vulnerabilities, an administrative user must also map vulnerabilities for the applications in the Firewall Management Center configuration.
- If you import application data and you want to use that data for impact correlation, you must map the vendor string for each application protocol to the corresponding Cisco application protocol definition.

Related Topics

Mapping Third-Party Products, on page 2598
Mapping Third-Party Product Fixes, on page 2599
Mapping Third-Party Vulnerabilities, on page 2600
Creating Custom Product Mappings, on page 2601

Third-Party Product Mappings

When you add data from third parties to the network map through the user input feature, you must map the vendor, product, and version names used by the third party to the Cisco product definitions. Mapping the products to Cisco definitions assigns vulnerabilities based on those definitions.

Similarly, if you are importing patch information from a third party, such as a patch management product, you must map the name for the fix to the appropriate vendor and product and the corresponding fix in the database.

Mapping Third-Party Products

If you import data from a third party, you must map the Cisco product to the third-party name to assign vulnerabilities and perform impact correlation using that data. Mapping the product associates Cisco vulnerability information with the third-party product name, which allows the system to perform impact correlation using that data.

If you import data using the host input import feature, you can also use the AddScanResult function to map third-party products to operating system and application vulnerabilities during the import.

For example, if you import data from a third party that lists Apache Tomcat as an application and you know it is version 6 of that product, you could add a third-party map where:

- **Vendor Name** is set to `Apache`.
- **Product Name** is set to `Tomcat`.
- **Apache** is chosen from the **Vendor** drop-down list.
- **Tomcat** is chosen from the **Product** drop-down list.
- **6** is chosen from the **Version** drop-down list

This mapping would cause any vulnerabilities for Apache Tomcat 6 to be assigned to hosts with an application listing for Apache Tomcat.

Note that for versionless or vendorless applications, you must map vulnerabilities for the application types in the Secure Firewall Management Center configuration. Although many clients have associated vulnerabilities, and clients are used for impact assessment, you cannot import and map third-party client vulnerabilities.

Tip If you have already created a third-party mapping on another Secure Firewall Management Center, you can export it and then import it onto this Firewall Management Center. You can then edit the imported mapping to suit your needs.

Procedure

Step 1 Choose **Policies** > **Application Detectors**.

Step 2 Click **User Third-Party Mappings**.

Step 3 You have two choices:

- Create — To create a new map set, click **Create Product Map Set**.
- Edit — To edit an existing map set, click **Edit** () next to the map set you want to modify. If **View** () appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Enter a **Mapping Set Name**.

Step 5 Enter a **Description**.

Step 6 You have two choices:

- Create — To map a third-party product, click **Add Product Map**.
- Edit — To edit an existing third-party product map, **Edit** () next to the map set you want to modify. If **View** () appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 7 Enter the **Vendor String** used by the third-party product.

Step 8 Enter the **Product String** used by the third-party product.

Step 9 Enter the **Version String** used by the third-party product.

Step 10 In the Product Mappings section, choose the operating system, product, and versions you want to use for vulnerability mapping from the **Vendor**, **Product**, **Major Version**, **Minor Version**, **Revision Version**, **Build**, **Patch**, and **Extension** fields.

Example:

If you want a host running a product whose name consists of third-party strings to use the vulnerabilities from Red Hat Linux 9, choose **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.

Step 11 Click **Save**.

Mapping Third-Party Product Fixes

If you map a fix name to a particular set of fixes in the database, you can then import data from a third-party patch management application and apply the fix to a set of hosts. When the fix name is imported to a host, the system marks all vulnerabilities addressed by the fix as invalid for that host.

Procedure

Step 1 Choose **Policies** > **Application Detectors**.

Step 2 Click **User Third-Party Mappings**.

Step 3 You have two choices:

- Create — To create a new map set, click **Create Product Map Set**.

- Edit — To edit an existing map set, click **Edit** (✎) next to the map set you want to modify. If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Enter a **Mapping Set Name**.

Step 5 Enter a **Description**.

Step 6 You have two choices:

- Create — To map a third-party product, click **Add Fix Map.**
- Edit — To edit an existing third-party product map, click **Edit** (✎) next to it. If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 7 Enter the name of the fix you want to map in the **Third-Party Fix Name** field.

Step 8 In the **Product Mappings** section, choose the operating system, product, and versions you want to use for fix mapping from the following fields:

- **Vendor**
- **Product**
- **Major Version**
- **Minor Version**
- **Revision Version**
- **Build**
- **Patch**
- **Extension**

Example:

If you want your mapping to assign the fixes from Red Hat Linux 9 to hosts where the patch is applied, choose **Redhat, Inc.** as the vendor, **Redhat Linux** as the product, and **9** as the version.

Step 9 Click **Save** to save the fix map.

Mapping Third-Party Vulnerabilities

To add vulnerability information from a third party to the VDB, you must map the third-party identification string for each imported vulnerability to any existing SVID, Bugtraq, or SID. After you create a mapping for the vulnerability, the mapping works for all vulnerabilities imported to hosts in the network map and allows impact correlation for those vulnerabilities.

You must enable impact correlation for third-party vulnerabilities to allow correlation to occur. For versionless or vendorless applications, you must also map vulnerabilities for the application types in the Secure Firewall Management Center configuration.

Although many clients have associated vulnerabilities, and clients are used for impact assessment, you cannot use third-party client vulnerabilities for impact assessment.

Tip If you have already created a third-party mapping on another Secure Firewall Management Center, you can export it and then import it onto this Firewall Management Center. You can then edit the imported mapping to suit your needs.

Procedure

Step 1 Choose **Policies** > **Application Detectors**.

Step 2 Click **User Third-Party Mappings**.

Step 3 You have two choices:

- Create — To create a new vulnerability set, click **Create Vulnerability Map Set**.
- Edit — To edit an existing vulnerability set, click **Edit** ( ) next to the vulnerability set. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Click **Add Vulnerability Map**.

Step 5 Enter the third-party identification for the vulnerability in the **Vulnerability ID** field.

Step 6 Enter a **Vulnerability Description**.

Step 7 Optionally:

- Enter a Snort ID in the **Snort Vulnerability ID Mappings** field.
- Enter a legacy vulnerability ID in the **SVID Mappings** field.
- Enter a Bugtraq identification number in the **Bugtraq Vulnerability ID Mappings** field.

Step 8 Click **Add**.

Related Topics

Enabling Network Discovery Vulnerability Impact Assessment, on page 2657

Custom Product Mappings

You can use product mappings to ensure that servers input by a third party are associated with the appropriate Cisco definitions. After you define and activate the product mapping, all servers or clients on monitored hosts that have the mapped vendor strings use the custom product mappings. For this reason, you may want to map vulnerabilities for all servers in the network map with a particular vendor string instead of explicitly setting the vendor, product, and version for the server.

Creating Custom Product Mappings

If the system cannot map a server to a vendor and product in the VDB, you can manually create the mapping. When you activate a custom product mapping, the system maps vulnerabilities for the specified vendor and product to all servers in the network map where that vendor string occurs.

Note Custom product mappings apply to all occurrences of an application protocol, regardless of the source of the application data (such as Nmap, the host input feature, or the system itself). However, if third-party vulnerability mappings for data imported using the host input feature conflicts with the mappings you set through a custom product mapping, the third-party vulnerability mapping overrides the custom product mapping and uses the third-party vulnerability mapping settings when the input occurs.

You create lists of product mappings and then enable or disable use of several mappings at once by activating or deactivating each list. When you specify a vendor to map to, the system updates the list of products to include only those made by that vendor.

After you create a custom product mapping, you must activate the custom product mapping list. After you activate a list of custom product mappings, the system updates all servers with occurrences of the specified vendor strings. For data imported through the host input feature, vulnerabilities update unless you have already explicitly set the product mappings for this server.

If, for example, your company modifies the banner for your Apache Tomcat web servers to read `Internal Web Server`, you can map the vendor string `Internal Web Server` to the vendor **Apache** and the product **Tomcat**, then activate the list containing that mapping, all hosts where a server labeled `Internal Web Server` occurs have the vulnerabilities for Apache Tomcat in the database.

Tip You can use this feature to map vulnerabilities to local intrusion rules by mapping the SID for the rule to another vulnerability.

Procedure

Step 1 Choose **Policies** > **Application Detectors**.

Step 2 Click **Custom Product Mappings**

Step 3 Click **Create Custom Product Mapping List**.

Step 4 Enter a **Custom Product Mapping List Name**.

Step 5 Click **Add Vendor String**.

Step 6 In the **Vendor String** field, enter the vendor string that identifies the applications that should map to the chosen vendor and product values.

Step 7 Choose the vendor you want to map to from the **Vendor** drop-down list.

Step 8 Choose the product you want to map to from the **Product** drop-down list.

Step 9 Click **Add** to add the mapped vendor string to the list.

Step 10 Optionally, repeat steps 4 to 8 as needed to add additional vendor string mappings to the list.

Step 11 Click **Save**.

What to do next

- Activate the custom product mapping list. For more information, see Activating and Deactivating Custom Product Mappings, on page 2603.

Editing Custom Product Mapping Lists

You can modify existing custom product mapping lists by adding or removing vendor strings or changing the list name.

Procedure

Step 1 Choose **Policies** > **Application Detectors**.

Step 2 Click **Custom Product Mappings**.

Step 3 Click **Edit** (✎) next to the product mapping list you want to edit.

If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Make changes to the list as described in Creating Custom Product Mappings, on page 2601.

Step 5 When you finish, click **Save**.

Activating and Deactivating Custom Product Mappings

You can enable or disable use of an entire list of custom product mappings at once. After you activate a custom product mapping list, each mapping on that list applies to all applications with the specified vendor string, whether detected by managed devices or imported through the host input feature.

Procedure

Step 1 Choose **Policies** > **Application Detectors**.

Step 2 Click **Custom Product Mappings**.

Step 3 Click the slider next to the custom product mapping list to activate or deactivate it.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Configuring the Host Input Client

The host input feature allows you to update the Firewall Management Center's network map from a client program running on another appliance. For example, you can add or delete hosts from the network map, or update the host OS and service information. For more information, see *Firepower System Host Input API Guide*.

Before you can run a remote client, you must add the client to the Firewall Management Center’s peers database from the Host Input Client page. You must also copy the authentication certificate generated by the Firewall Management Center to the client. After completing these steps the client can connect to the Firewall Management Center.

In a multidomain deployment, you can create a client in any domain. The authentication certificate allows the client to submit network map updates for any leaf domains associated with the client certificate's domain. If you create a certificate for an ancestor domain (or if your certificate domain later becomes an ancestor domain after adding descendant domains), any clients using that certificate must specify a target leaf domain with every transaction, as described in the *Firepower System Host Input API Guide*.

The Host Input Client shows only clients associated the current domain, so if you want to download or revoke a certificate, switch to the domain where the client was created.

This connection uses TLS 1.2.

Procedure

Step 1 Choose **Integration** > **Other Integrations**.

Step 2 Click **Host Input Client**.

Step 3 Click **Create Client**.

Step 4 In the **Hostname** field, enter the host name or IP address of the host running the host input client.

Note
If you have not configured DNS resolution, use an IP address.

Step 5 If you want to encrypt the certificate file, enter a password in the **Password** field.

Step 6 Click **Save**.
The host input service allows the host to access port 8307 on the Firewall Management Center and creates an authentication certificate to use during client-server authentication.

Step 7 Click **Download** ( ) next to the certificate file.

Step 8 Save the certificate file to the directory used by your client for SSL/TLS authentication.

Step 9 To revoke access for a client, click **Delete** ( ) next to the host you want to remove.

Nmap Scanning

The system builds network maps through passive analysis of traffic on your network. Information obtained through this passive analysis can occasionally be incomplete, depending on system conditions. However, you can actively scan a host to obtain complete information. For example, if a host has a server running on an open port but the server has not received or sent traffic during the time that the system has been monitoring your network, the system does not add information about that server to the network map. If you directly scan that host using an active scanner, however, you can detect the presence of the server.

The system integrates with Nmap™, an open source active scanner for network exploration and security auditing.

When you scan a host using Nmap, the system:

- Adds servers on previously undetected open ports to the Servers list in the host profile for that host. The host profile lists any servers detected on filtered or closed TCP ports or on UDP ports in the Scan Results section. By default, Nmap scans more than 1660 TCP ports.

  If the system recognizes a server identified in an Nmap scan and has a corresponding server definition, the system maps the names Nmap uses for servers to the corresponding Cisco server definitions.

- Compares the results of the scan to over 1500 known operating system fingerprints to determine the operating system and assigns scores to each. The operating system assigned to the host is the operating system fingerprint with the highest score.

The system maps Nmap operating system names to Cisco operating system definitions.

- Assigns vulnerabilities to the host for the added servers and operating systems.

Note:

- A host must exist in the network map before Nmap can append its results to the host profile.
- If the host is deleted from the network map, any Nmap scan results for that host are discarded.

Tip Some scanning options (such as portscans) may place a significant load on networks with low bandwidths. Schedule scans like these to run during periods of low network use.

For more information on the underlying Nmap technology used to scan, refer to the Nmap documentation at http://insecure.org/.

Nmap Remediation Options

You define the settings for an Nmap scan by creating an Nmap remediation. An Nmap remediation can be used as a response in a correlation policy, run on demand, or scheduled to run at a specific time.

Note that Nmap-supplied server and operating system data remain static until you run another Nmap scan. If you plan to scan a host for operating system and server data using Nmap, you may want to set up regularly scheduled scans to keep any Nmap-supplied operating system and server data up-to-date.

The following table explains the options configurable in Nmap remediations.

Table 139: Nmap Remediation Options

Option	Description	Corresponding Nmap Option
Scan Which Address(es) From Event?	When you use an Nmap scan as a response to a correlation rule, select one of the following options to control which address in the event is scanned, that of the source host, the destination host, or both: • **Scan Source and Destination Addresses** scans the hosts represented by the source IP address and the destination IP address in the event. • **Scan Source Address Only** scans the host represented by the event's source IP address. • **Scan Destination Address Only** scans the host represented by the event's destination IP address.	N/A

Option	Description	Corresponding Nmap Option
Scan Types	Select how Nmap scans ports: • The **TCP Syn** scan connects quickly to thousand of ports without using a complete TCP handshake. This options allows you to scan quickly in stealth mode on hosts where the `admin` account has raw packet access or where IPv6 is not running, by initiating TCP connections but not completing them. If a host acknowledges the Syn packet sent in a TCP Syn scan, Nmap resets the connection. • The **TCP Connect** scan uses the `connect()` system call to open connections through the operating system on the host. You can use the TCP Connect scan if the `admin` user on the Firewall Management Center or managed device does not have raw packet privileges on a host or you are scanning IPv6 networks. In other words, use this option in situations where the TCP Syn scan cannot be used. • The **TCP ACK** scan sends an ACK packet to check whether ports are filtered or unfiltered. • The **TCP Window** scan works in the same way as a TCP ACK scan but can also determine whether a port is open or closed. • The **TCP Maimon** scan identifies BSD-derived systems using a FIN/ACK probe.	**TCP Syn**: `-sS` **TCP Connect**: `-sT` **TCP ACK**: `-sA` **TCP Window**: `-sW` **TCP Maimon**: `-sM`
Scan for UDP ports	Enable to scan UDP ports in addition to TCP ports. Note that scanning UDP ports may be time-consuming, so avoid using this option if you want to scan quickly.	`-sU`
Use Port From Event	If you plan to use the remediation as a response in a correlation policy, enable to cause the remediation to scan only the port specified in the event that triggers the correlation response. • Select **On** to scan the port in the correlation event, rather than the ports you specify during Nmap remediation configuration. If you scan the port in the correlation event, note that the remediation scans the port on the IP addresses that you specify during Nmap remediation configuration. These ports are also added to the remediation's dynamic scan target. • Select **Off** to scan only the ports you specify Nmap remediation configuration. You can also control whether Nmap collects information about operating system and server information. Enable the **Use Port From Event** option to scan the port associated with the new server.	N/A

Option	Description	Corresponding Nmap Option
Scan from reporting detection engine	Enable to scan a host from the appliance where the detection engine that reported the host resides. • To scan from the appliance running the reporting detection engine, select **On**. • To scan from the appliance configured in the remediation, select **Off**.	N/A
Fast Port Scan	Enable to scan only the TCP ports listed in the `nmap-services` file located in the `/var/sf/nmap/share/nmap/nmap-services` directory on the device that does the scanning, ignoring other port settings. Note that you cannot use this option with the **Port Ranges and Scan Order** option. • To scan only the ports listed in the `nmap-services` file located in the `/var/sf/nmap/share/nmap/nmap-services` directory on the device that does the scanning, ignoring other port settings, select **On**. • To scan all TCP ports, select **Off**.	`-F`
Port Ranges and Scan Order	Set the specific ports you want to scan, using Nmap port specification syntax, and the order you want to scan them. Note that you cannot use this option with the **Fast Port Scan** option.	`-p`
Probe open ports for vendor and version information	Enable to detect server vendor and version information. If you probe open ports for server vendor and version information, Nmap obtains server data that it uses to identify servers. It then replaces the Cisco server data for that server. • Select **On** to scan open ports on the host for server information to identify server vendors and versions. • Select **Off** to continue using Cisco server information for the host.	`-sV`
Service Version Intensity	Select the intensity of Nmap probes for service versions. • To use more probes for higher accuracy with a longer scan, select a higher number. • To use fewer probes for less accuracy with a faster scan, select a lower number.	`--version-intensity` <intensity>
Detect Operating System	Enable to detect operating system information for the host. If you configure detection of the operating system for a host, Nmap scans the host and uses the results to create a rating for each operating system that reflects the likelihood that the operating system is running on the host. • Select **On** to scan the host for information to identify the operating system. • Select **Off** to continue using Cisco operating system information for the host.	`-o`

Option	Description	Corresponding Nmap Option
Treat All Hosts As Online	Enable to skip the host discovery process and run a port scan on every host in the target range. Note that when you enable this option, Nmap ignores settings for **Host Discovery Method** and **Host Discovery Port List**. • To skip the host discovery process and run a port scan on every host in the target range, select **On**. • To perform host discovery using the settings for **Host Discovery Method** and **Host Discovery Port List** and skip the port scan on any host that is not available, select **Off**.	`-PN`
Host Discovery Method	Select to perform host discovery for all hosts in the target range, over the ports listed in the **Host Discovery Port Lis**t, or if no ports are listed, over the default ports for that host discovery method. Note that if you also enabled **Treat All Hosts As Online**, however, the **Host Discovery Method** option has no effect and host discovery is not performed. Select the method to be used when Nmap tests to see if a host is present and available: • The **TCP SYN** option sends an empty TCP packet with the SYN flag set and recognizes the host as available if a response is received. TCP SYN scans port 80 by default. Note that TCP SYN scans are less likely to be blocked by a firewall with stateful firewall rules. • The **TCP ACK** option sends an empty TCP packet with the ACK flag set and recognizes the host as available if a response is received. TCP ACK also scans port 80 by default. Note that TCP ACK scans are less likely to be blocked by a firewall with stateless firewall rules. • The **UDP** option sends a UDP packet and assumes host availability if a port unreachable response comes back from a closed port. UDP scans port 40125 by default.	**TCP SYN**: `-PS` **TCP ACK**: `-PA` **UDP**: `-PU`
Host Discovery Port List	Specify a customized list of ports, separated by commas, that you want to scan when doing host discovery.	port list for host discovery method
Default NSE Scripts	Enable to run the default set of Nmap scripts for host discovery and server and operating system and vulnerability detection. See https://nmap.org/nsedoc/categories/default.html for the list of default scripts. • To run the default set of Nmap scripts, select **On**. • To skip the default set of Nmap scripts, select **Off**.	`-sC`

Option	Description	Corresponding Nmap Option
Timing Template	Select the timing of the scan process; the higher the number you select, the faster and less comprehensive the scan.	**0**: `T0` (paranoid) **1**: `T1` (sneaky) **2**: `T2` (polite) **3**: `T3` (normal) **4**: `T4` (aggressive) **5**: `T5` (insane)

Nmap Scanning Guidelines

While active scanning can obtain valuable information, overuse of a tool such as Nmap may overload your network resources or even crash important hosts. When using any active scanner, you should create a scanning strategy following these guidelines to make sure that you are scanning only the hosts and ports that you need to scan.

Selecting Appropriate Scan Targets

When you configure Nmap, you can create scan targets that identify which hosts you want to scan. A scan target includes a single IP address, a CIDR block or octet range of IP addresses, an IP address range, or a list of IP addresses or ranges to scan, as well as the ports on the host or hosts.

You can specify targets in the following ways:

- For IPv6 hosts:
  - an exact IP address (for example, `2001:DB8:1::178:ABCD`)
- For IPv4 hosts:
  - an exact IP address (for example, `192.168.1.101`) or a list of IP addresses separated by commas or spaces
  - an IP address block using CIDR notation (for example, `192.168.1.0/24` scans the 254 hosts between 192.168.1.1 and 192.168.1.254, inclusive).
  - an IP address range using octet range addressing (for example, `192.168.0-255.1-254` scans all addresses in the 192.168.x.x range, except those that end in .0 and or .255)
  - an IP address range using hyphenation (for example, `192.168.1.1 - 192.168.1.5` scans the six hosts between 192.168.1.1 and 192.168.1.5, inclusive)
  - a list of addresses or ranges separated by commas or spaces (for example, for example, `192.168.1.0/24, 194.168.1.0/24` scans the 254 hosts between 192.168.1.1 and 192.168.1.254, inclusive and the 254 hosts between 194.168.1.1 and 194.168.1.254, inclusive)

Ideal scan targets for Nmap scans include hosts with operating systems that the system is unable to identify, hosts with unidentified servers, or hosts recently detected on your network. Remember that Nmap results cannot be added to the network map for hosts that do not already exist in the network map.

Caution

- Nmap-supplied server and operating system data remains static until you run another Nmap scan. If you plan to scan a host using Nmap, regularly schedule scans.
- If a host is deleted from the network map, any Nmap scan results are discarded.
- Make sure you have permission to scan your targets. Using Nmap to scan hosts that do not belong to you or your company may be illegal.

Selecting Appropriate Ports to Scan

For each scan target you configure, you can select the ports you want to scan. You can designate individual port numbers, port ranges, or a series of port numbers and port ranges to identify the exact set of ports that should be scanned on each target.

By default, Nmap scans TCP ports 1 through 1024. If you plan to use the remediation as a response in a correlation policy, you can cause the remediation to scan only the port specified in the event that triggers the correlation response. If you run the remediation on demand or as a scheduled task, or if you do not use the port from the event, you can use other port options to determine which ports are scanned. You can choose to scan only the TCP ports listed in the `nmap-services` file, ignoring other port settings. You can also scan UDP ports in addition to TCP ports. Note that scanning for UDP ports may be time-consuming, so avoid using that option if you want to scan quickly. To select the specific ports or range of ports to scan, use Nmap port specification syntax to identify ports.

Setting Host Discovery Options

You can decide whether to perform host discovery before starting a port scan for a host, or you can assume that all the hosts you plan to scan are online. If you choose not to treat all hosts as online, you can choose what method of host discovery to use and, if needed, customize the list of ports scanned during host discovery. Host discovery does not probe the ports listed for operating system or server information; it uses the response over a particular port only to determine whether a host is active and available. If you perform host discovery and a host is not available, Nmap does not scan ports on that host.

Example: Using Nmap to Resolve Unknown Operating Systems

This example walks through an Nmap configuration designed to resolve unknown operating systems. For a complete look at Nmap configuration, see Managing Nmap Scanning, on page 2612.

If the system cannot determine the operating system on a host on your network, you can use Nmap to actively scan the host. Nmap uses the information it obtains from the scan to rate the possible operating systems. It then uses the operating system that has the highest rating as the host operating system identification.

Using Nmap to challenge new hosts for operating system and server information deactivates the system's monitoring of that data for scanned hosts. If you use Nmap to discover host and server operating system for hosts the system marks as having unknown operating systems, you may be able to identify groups of hosts that are similar. You can then create a custom fingerprint based on one of them to cause the system to associate the fingerprint with the operating system you know is running on the host based on the Nmap scan. Whenever possible, create a custom fingerprint rather than inputting static data through a third-party source like Nmap because the custom fingerprint allows the system to continue to monitor the host operating system and update it as needed.

In this example, you would:

1. Configure a scan instance as described in Adding an Nmap Scan Instance, on page 2613.
2. Create an Nmap remediation using the following settings:
   - Enable **Use Port From Event** to scan the port associated with the new server.
   - Enable **Detect Operating System** to detect operating system information for the host.
   - Enable **Probe open ports for vendor and version information** to detect server vendor and version information.
   - Enable **Treat All Hosts as Online**, because you know the host exists.
3. Create a correlation rule that triggers when the system detects a host with an unknown operating system. The rule should trigger when **a discovery event occurs** and **the OS information for a host has changed** and it meets the following conditions: **OS Name is unknown**.
4. Create a correlation policy that contains the correlation rule.
5. n the correlation policy, add the Nmap remediation you created in step 2 as a response to the rule you created in step 3.
6. Activate the correlation policy.
7. Purge the hosts on the network map to force network discovery to restart and rebuild the network map.
8. After a day or two, search for events generated by the correlation policy. Analyze the Nmap results for the operating systems detected on the hosts to see if there is a particular host configuration on your network that the system does not recognize.
9. If you find hosts with unknown operating systems whose Nmap results are identical, create a custom fingerprint for one of those hosts and use it to identify similar hosts in the future.

Related Topics

Creating an Nmap Remediation, on page 2616
Nmap Scan Results, on page 2619
Creating a Custom Fingerprint for Clients, on page 2592

Example: Using Nmap to Respond to New Hosts

This example walks through an Nmap configuration designed to respond to new hosts. For a complete look at Nmap configuration, see Managing Nmap Scanning, on page 2612.

When the system detects a new host in a subnet where intrusions may be likely, you may want to scan that host to make sure you have accurate vulnerability information for it.

You can accomplish this by creating and activating a correlation policy that detects when a new host appears in this subnet, and that launches a remediation that performs an Nmap scan on the host.

To do this, you would:

1. Configure a scan instance as described in Adding an Nmap Scan Instance, on page 2613.
2. Create an Nmap remediation using the following settings:
   - Enable **Use Port From Event** to scan the port associated with the new server.
   - Enable **Detect Operating System** to detect operating system information for the host.

- Enable **Probe open ports for vendor and version information** to detect server vendor and version information.
- Enable **Treat All Hosts as Online**, because you know the host exists.

3. Create a correlation rule that triggers when the system detects a new host on a specific subnet. The rule should trigger when **a discovery event occurs** and **a new host is detected**.
4. Create a correlation policy that contains the correlation rule.
5. In the correlation policy, add the Nmap remediation you created in step 2 as a response to the rule you created in step 3.
6. Activate the correlation policy.
7. When you are notified of a new host, check the host profile to see the results of the Nmap scan and address any vulnerabilities that apply to the host.

After you activate the policy, you can periodically check the remediation status view (**Analysis** > **Correlation** > **Status**) to see when the remediation launched. The remediation's dynamic scan target should include the IP addresses of the hosts it scanned as a result of the server detection. Check the host profile for those hosts to see if there are vulnerabilities that need to be addressed for the host, based on the operating system and servers detected by Nmap.

Caution If you have a large or dynamic network, detection of a new host may be too frequent an occurrence to respond to using a scan. To prevent resource overload, avoid using Nmap scans as a response to events that occur frequently. In addition, note that using Nmap to challenge new hosts for operating system and server information deactivates Cisco monitoring of that data for scanned hosts.

Related Topics

Creating an Nmap Remediation, on page 2616

Managing Nmap Scanning

To use Nmap scanning, you must, at minimum, configure an Nmap scan instance and an Nmap remediation. Configuring an Nmap scan target is optional.

Procedure

Step 1 Configure the Nmap scan:

- Add an Nmap scan instance as described in Adding an Nmap Scan Instance, on page 2613.
- Create an Nmap remediation as described in Creating an Nmap Remediation, on page 2616.
- Optionally, add an Nmap scan target as described in Adding an Nmap Scan Target, on page 2615.

Step 2 Run the Nmap scan:

- Run an on-demand Nmap scan as described in Running an On-Demand Nmap Scan, on page 2619.
- Configure automatic Nmap scans as described in *Nmap Scan Automation* in the Cisco Secure Firewall Management Center Administration Guide.

- Schedule automatic Nmap scans as described in *Scheduling an Nmap Scan* in the Cisco Secure Firewall Management Center Administration Guide.

What to do next

- Monitor the Nmap scan in progress by viewing the related task; see *Viewing Task Messages* in the Cisco Secure Firewall Management Center Administration Guide.
- Optionally, refine the scan:
  - Edit an Nmap scan instance as described in Editing an Nmap Scan Instance, on page 2614.
  - Edit an Nmap scan target as described in Editing an Nmap Scan Target, on page 2616.
  - Edit an Nmap remediation as described in Editing an Nmap Remediation, on page 2618.

Adding an Nmap Scan Instance

You can set up a separate scan instance for each Nmap module that you want to use to scan your network for vulnerabilities. You can set up scan instances for the local Nmap module on the Secure Firewall Management Center and for any devices you want to use to run scans remotely. The results of each scan are always stored on the Firewall Management Center where you configure the scan, even if you run the scan from a remote device. To prevent accidental or malicious scanning of mission-critical hosts, you can create a blacklist for the instance to indicate the hosts that should never be scanned with the instance.

You cannot add a scan instance with the same name as any existing scan instance.

Procedure

Step 1 Access the list of Nmap scan instances using either of the following methods:

- Choose **Policies** > **Actions** > **Instances**.
- Choose **Policies** > **Actions** > **Scanners**.

Step 2 Add the remediation:

- If you accessed the list via the first method above, locate the Add a New Instance section, choose the Nmap Remediation module from the drop-down list, and click **Add**.
- If you accessed the list via the second method above, click **Add Nmap Instance**.

Step 3 Enter an **Instance Name**.

Step 4 Enter a **Description**.

Step 5 Optionally, in the **Exempted hosts** field, specify any hosts or networks that should *never* be scanned with this scan instance, using the following syntax:

- For IPv6 hosts, an exact IP address (for example, `2001:DB8::fedd:eeff`)
- For IPv4 hosts, an exact IP address (for example, `192.168.1.101`) or an IP address block using CIDR notation (for example, `192.168.1.0/24` scans the 254 hosts between `192.168.1.1` and `192.168.1.254`, inclusive)

- Note that you cannot use an exclamation mark (!) to negate an address value.

Note
If you specifically target a scan to a host that is in a blacklisted network, that scan will not run.

Step 6 Optionally, to run the scan from a remote device instead of the Firewall Management Center, specify the IP address or name of the device as it appears in the Information page for the device in the Firewall Management Center web interface, in the **Remote Device Name** field.

Step 7 Click **Create**.
When the system is done creating the instance, it displays it in edit mode.

Step 8 Optionally, add an Nmap remediation to the instance. To do so, locate the Configured Remediations section of the instance, click **Add**, and create a remediation as described in Creating an Nmap Remediation, on page 2616.

Step 9 Click **Cancel** to return to the list of instances.

Note
If you accessed the list of Nmap scan instances via the **Scanners** option, the system does not display the instance you added unless you also added a remediation to the instance. To view any instances to which you have not yet added remedations, use the **Instances** menu option to access the list.

Editing an Nmap Scan Instance

When you edit a scan instance, you can view, add, and delete remediations associated with the instance. Delete an Nmap scan instance when you no longer want to use the Nmap module profiled in the instance. Note that when you delete the scan instance, you also delete any remediations that use that instance.

Procedure

Step 1 Access the list of Nmap scan instances using either of the following methods:

- Choose **Policies** > **Actions** > **Instances**.
- Choose **Policies** > **Actions** > **Scanners**.

Step 2 Click **View** ( ) next to the instance you want to edit.

Step 3 Make changes to the scan instance settings as described in Adding an Nmap Scan Instance, on page 2613.

Step 4 Click **Save**.

Step 5 Click **Done**.

What to do next

- Optionally, add a new remediation to the scan instance; see Creating an Nmap Remediation, on page 2616
- Optionally, edit a remediation associated with the instance; see Editing an Nmap Remediation, on page 2618.

- Optionally, delete a remediation associated with the instance; see Running an On-Demand Nmap Scan, on page 2619.
- Optionally, delete the scan instance by clicking **Delete** ( ) next to it.

Adding an Nmap Scan Target

When you configure an Nmap module, you can create and save scan targets that identify the hosts and ports you want to target when you perform an on-demand or a scheduled scan, so that you do not have to construct a new scan target every time. A scan target includes a single IP address or a block of IP addresses to scan, as well as the ports on the host or hosts. For Nmap targets, you can also use Nmap octet range addressing or IP address ranges. For more information on Nmap octet range addressing, refer to the Nmap documentation at http://insecure.org.

Note:

- Scans for scan targets containing a large number of hosts can take an extended period of time. As a workaround, scan fewer hosts at a time.
- Nmap-supplied server and operating system data remains static until you run another Nmap scan. If you plan to scan a host using Nmap, regularly schedule scans. If a host is deleted from the network map, any Nmap scan results are discarded.

Procedure

Step 1 Choose **Policies** > **Actions** > **Scanners**.

Step 2 On the toolbar, click **Targets**.

Step 3 Click **Create Scan Target**.

Step 4 In the **Name** field, enter the name you want to use for this scan target.

Step 5 In the **IP Range** text box, specify the host or hosts you want to scan using the syntax described in Nmap Scanning Guidelines, on page 2609.

Note
If you use a comma in a list of IP addresses or ranges in a scan target, the comma converts to a space when you save the target.

Step 6 In the **Ports** field, specify the ports you want to scan.

You can enter any of the following, using values from 1 to 65535:

- a port number
- a list of ports separated by commas
- a range of port numbers separated by a dash
- ranges of port numbers separated by dashes, separated by commas

Step 7 Click **Save**.

Editing an Nmap Scan Target

Tip You might want to edit a remediation's dynamic scan target if you do not want to use the remediation to scan a specific IP address, but the IP address was added to the target because the host was involved in a correlation policy violation that launched the remediation.

Delete a scan target if you no longer want to scan the hosts listed in it.

Procedure

Step 1 Choose **Policies** > **Actions** > **Scanners**.

Step 2 On the toolbar, click **Targets**.

Step 3 Click **Edit** () next to the scan target you want to edit.

If **View** () appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 4 Make modifications as necessary. For more information, see Adding an Nmap Scan Target, on page 2615.

Step 5 Click **Save**.

Step 6 Optionally, delete the scan target by clicking **Delete** () next to it.

Creating an Nmap Remediation

An Nmap remediation can only be created by adding it to an existing Nmap scan instance. The remediation defines the settings for the scan. It can be used as a response in a correlation policy, run on demand, or run as a scheduled task at a specific time.

Nmap-supplied server and operating system data remains static until you run another Nmap scan. If you plan to scan a host using Nmap, regularly schedule scans. If a host is deleted from the network map, any Nmap scan results are discarded.

For general information about Nmap functionality, refer to the Nmap documentation at http://insecure.org.

Before you begin

- Add an Nmap scan instanceas described in Adding an Nmap Scan Instance, on page 2613 .

Procedure

Step 1 Choose **Policies** > **Actions** > **Instances**.

Step 2 Click **View** () next to the instance to which you want to add the remediation.

Step 3 In the Configured Remediations section, click **Add**.

Step 4 Enter a **Remediation Name**.

Step 5 Enter a **Description**.

Step 6 If you plan to use this remediation in response to a correlation rule that triggers on an intrusion event, a connection event, or a user event, configure the **Scan Which Address(es) From Event?** option.

Tip
If you plan to use this remediation in response to a correlation rule that triggers on a discovery event or a host input event, by default the remediation scans the IP address of the host involved in the event; you do not need to configure this option.

Note
Do **not** assign an Nmap remediation as a response to a correlation rule that triggers on a traffic profile change.

Step 7 Configure the **Scan Type** option.

Step 8 Optionally, to scan UDP ports in addition to TCP ports, choose **On** for the **Scan for UDP ports** option.

Tip
A UDP portscan takes more time than a TCP portscan. To speed up your scans, leave this option disabled.

Step 9 If you plan to use this remediation in response to correlation policy violations, configure the **Use Port From Event** option.

Step 10 If you plan to use this remediation in response to correlation policy violations and want to run the scan using the appliance running the detection engine that detected the event, configure the **Scan from reporting detection engine** option.

Step 11 Configure the **Fast Port Scan** option.

Step 12 In the **Port Ranges and Scan Order** field, enter the ports you want to scan by default, using Nmap port specification syntax, in the order you want to scan those ports.

Use the following format:

- Specify values from 1 to 65535.
- Separate ports using commas or spaces.
- Use a hyphen to indicate a port range.
- When scanning for both TCP and UDP ports, preface the list of TCP ports you want to scan with a T and the list of UDP ports with a U.

Note
The **Use Port From Event** option overrides this setting when the remediation is launched in response to a correlation policy violation, as described in step 8.

Example:

To scan ports 53 and 111 for UDP traffic, then scan ports 21-25 for TCP traffic, enter `U:53,111,T:21-25`.

Step 13 To probe open ports for server vendor and version information, configure **Probe open ports for vendor and version information**.

Step 14 If you choose to probe open ports, set the number of probes used by choosing a number from the **Service Version Intensity** drop-down list.

Step 15 To scan for operating system information, configure **Detect Operating System** settings.

Step 16 To determine whether host discovery occurs and whether port scans are only run against available hosts, configure **Treat All Hosts As Online**.

Step 17 To set the method you want Nmap to use when it tests for host availability, choose a method from the **Host Discovery Method** drop-down list.

Step 18 If you want to scan a custom list of ports during host discovery, enter a list of ports appropriate for the host discovery method you chose, separated by commas, in the **Host Discovery Port List** field.

Step 19 Configure the **Default NSE Scripts** option to control whether to use the default set of Nmap scripts for host discovery and server, operating system, and vulnerability discovery.

Tip
See http://nmap.org/nsedoc/categories/default.html for the list of default scripts.

Step 20 To set the timing of the scan process, choose a timing template number from the **Timing Template** drop-down list.

Choose a higher number for a faster, less comprehensive scan and a lower number for a slower, more comprehensive scan.

Step 21 Click **Create**.
When the system is done creating the remediation, it displays it in edit mode.

Step 22 Click **Done** to return to the related instance.

Step 23 Click **Cancel** to return to the instance list.

Related Topics

Nmap Remediation Options, on page 2605

Editing an Nmap Remediation

Modifications you make to Nmap remediations do not affect scans in progress. The new settings take effect when the next scan starts. Delete an Nmap remediation if you no longer need it.

Procedure

Step 1 Access the list of Nmap scan instances using either of the following methods:

- Choose **Policies** > **Actions** > **Instances**.
- Choose **Policies** > **Actions** > **Scanners**.

Step 2 Access the remediation you want to edit:

- If you accessed the list via the first method above, click **View** ( ) next to the relevant instance, and then click it again next to the remediation you want to edit in the Configured Remediations section.
- If you accessed the list via the second method above, click **View** ( ) next to the remediation you want to edit.

Step 3 Make modifications as necessary as described in Creating an Nmap Remediation, on page 2616.

Step 4 Click **Save** if you want to save your changes, or **Done** if you want to exit without saving.

Step 5 Optionally, delete the remediation by clicking **Delete** ( ) next to it.

Running an On-Demand Nmap Scan

You can launch on-demand Nmap scans whenever needed. You can specify the target for an on-demand scan by entering the IP addresses and ports you want to scan or by choosing an existing scan target.

Nmap-supplied server and operating system data remains static until you run another Nmap scan. If you plan to scan a host using Nmap, regularly schedule scans. If a host is deleted from the network map, any Nmap scan results are discarded.

Before you begin

- Optionally, add an Nmap scan target; see Adding an Nmap Scan Target, on page 2615.

Procedure

Step 1 Choose **Policies** > **Actions** > **Scanners**.

Step 2 Next to the Nmap remediation you want to use to perform the scan, click **Scan** ( ).

Step 3 Optionally, to scan using a saved scan target, choose a target from the **Saved Targets** drop-down list, and click **Load**.

Step 4 In the **IP Range(s)** field, specify the IP address for hosts you want to scan or modify the loaded list.

Note:

- For hosts with IPv4 addresses, you can specify multiple IP addresses separated by commas or use CIDR notation. You can also negate IP addresses by preceding them with an exclamation point (!).
- For hosts with IPv6 addresses, use an exact IP address. Ranges are not supported.

Step 5 In the **Ports** field, specify the ports you want to scan or modify the loaded list.

You can enter a port number, a list of ports separated by commas, or a range of port numbers separated by a dash.

Step 6 Click **Scan Now**.

What to do next

- Optionally, monitor the task status; see *Viewing Task Messages* in the Cisco Secure Firewall Management Center Administration Guide.

Nmap Scan Results

You can monitor Nmap scans in progress, import results from scans previously performed through the system or results preformed outside the system, and view and analyze scan results.

You can view scan results that you create using the local Nmap module as a rendered page in a pop-up window. You can also download the Nmap results file in raw XML format.

You can also view operating system and server information detected by Nmap in host profiles and in the network map. If a scan of a host produces server information for servers on filtered or closed ports, or if a scan collects information that cannot be included in the operating system information or the servers section, the host profile includes those results in an Nmap Scan Results section.

Viewing Nmap Scan Results

When an Nmap scan is complete, you can view a table of scan results.

You can manipulate the results view depending on the information you are looking for. The page you see when you access scan results differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of scan results. You can also create a custom workflow that displays only the information that matches your specific needs.

You can download and view the Nmap results using the Nmap Version 1.01 DTD, available at http://insecure.org.

You can also clear scan results.

Procedure

Step 1 Choose **Policies** > **Actions** > **Scanners**.

Step 2 On the toolbar, click **Scan Results**.

Step 3 You have the following choices:

- Adjust the time range as described in *Event Time Constraints* in the Cisco Secure Firewall Management Center Administration Guide.
- To use a different workflow, including a custom workflow, click **(switch workflows)** by the workflow title.
- To view the scan results as a rendered page in a pop-up window, click **View** next to the scan job.
- To save a copy of the scan results file so that you can view the raw XML code in any text editor, click **Download** next to the scan job.
- To sort scan results, click the column title. Click the column title again to reverse the sort order.
- To constrain the columns that appear, click **Close** (✕) in the column heading that you want to hide. In the pop-up window that appears, click **Apply**.

  Tip
  To hide or show other columns, check or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, Click the expand arrow to expand the search constraints, then click the column name under **Disabled Columns**.

- To drill down to the next page in the workflow, see *Using Drill-Down Pages* in the Cisco Secure Firewall Management Center Administration Guide.
- To configure scan instances and remediations, click **Scanners** in the toolbar and see Managing Nmap Scanning, on page 2612.
- To navigate within and between workflow pages, see *Workflow Page Navigation Tools* in the Cisco Secure Firewall Management Center Administration Guide.
- To navigate to other event views to view associated events, choose the name of the event view you want to see from the **Jump to** drop-down list.

- To search for scan results, enter your search criteria in the appropriate fields.

Related Topics

Nmap Scan Results Fields, on page 2621

Nmap Scan Results Fields

When you run an Nmap scan, the Firewall Management Center collects the scan results in a database. The following table describes the fields in the scan results table that can be viewed and searched.

Table 140: Scan Results Fields

Field	Description
Start Time	The date and time that the scan that produced the results started.
End Time	The date and time that the scan that produced the results ended.
Target	The IP address (or host name, if DNS resolution is enabled) of the scan target for the scan that produced the results.
Scan Type	Either `Nmap` or the name of the third-party scanner to indicate the type of the scan that produced the results.
Scan Mode	The mode of the scan that produced the results: • `On Demand` — results from scans run on demand. • `Imported` — results from scans on a different system and imported onto the Firewall Management Center. • `Scheduled` — results from scans run as a scheduled task.
Results	The results of the scan.
Domain	The domain of the scan target. This field is only present in a multidomain deployment.

Importing Nmap Scan Results

You can import XML results files created by an Nmap scan performed outside of the system. You can also import XML results files that you previously downloaded from the system. To import Nmap scan results, the results file must be in XML format and adhere to the Nmap Version 1.01 DTD. For more information on creating Nmap results and on the Nmap DTD, refer to the Nmap documentation at http://insecure.org.

A host must already exist in the network map before Nmap can append its results to the host profile.

Procedure

Step 1 Choose **Policies** > **Actions** > **Scanners**.

Step 2 On the toolbar, click **Import Results**.

Step 3 Click **Browse** to navigate to the results file.

Step 4 After you return to the Import Results page, click **Import** to import the results.

History for Host Identity Sources

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Security improvement to the host input data feature	6.5	Any	TLS 1.2 is now used for communication between your Firewall Management Center and the host input client. The topic Configuring the Host Input Client, on page 2603has been updated with this information.

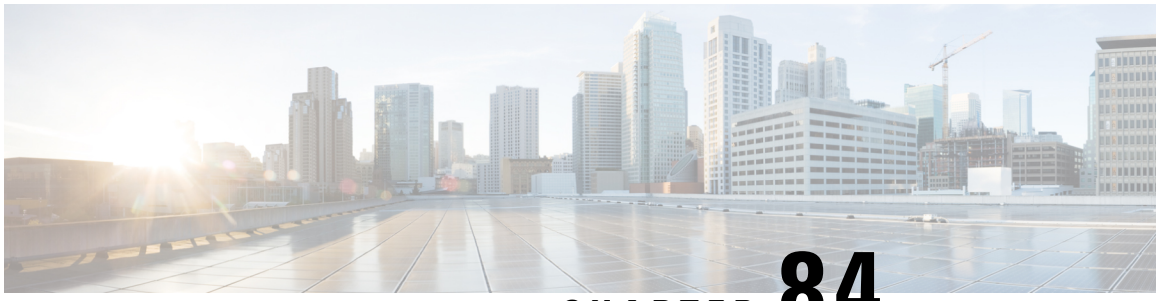

CHAPTER 84

Application Detection

The following topics describe application detection:

- Overview: Application Detection, on page 2623
- Requirements and Prerequisites for Application Detection, on page 2629
- Custom Application Detectors, on page 2629
- Viewing or Downloading Detector Details, on page 2638
- Sorting the Detector List, on page 2638
- Filtering the Detector List, on page 2638
- Navigating to Other Detector Pages, on page 2640
- Activating and Deactivating Detectors, on page 2640
- Editing Custom Application Detectors, on page 2641
- Deleting Detectors, on page 2642

Overview: Application Detection

When the system analyzes IP traffic, it attempts to identify the commonly used applications on your network. Application awareness is crucial to application control.

There are three types of applications that the system detects:

- *application protocols* such as HTTP and SSH, which represent communications between hosts
- *clients* such as web browsers and email clients, which represent software running on the host
- *web applications* such as MPEG video and Facebook, which represent the content or requested URL for HTTP traffic

The system identifies applications in your network traffic according to the characteristics specified in the detector. For example, the system can identify an application by an ASCII pattern in the packet header. In addition, Secure Socket Layers (SSL) protocol detectors use information from the secured session to identify the application from the session.

There are two sources of application detectors:

- *System-provided detectors* detect web applications, clients, and application protocols.

  The availability of system-provided detectors for applications (and operating systems) depends on the version of the system software and the version of the VDB you have installed. Release notes and advisories

contain information on new and updated detectors. You can also import individual detectors authored by Professional Services.

- *Custom application protocol detectors* are user-created and detect web applications, clients, and application protocols.

You can also detect application protocols through *implied application protocol detection*, which infers the existence of an application protocol based on the detection of a client.

The system identifies only those application protocols running on hosts in your monitored networks, as defined in the network discovery policy. For example, if an internal host accesses an FTP server on a remote site that you are not monitoring, the system does not identify the application protocol as FTP. On the other hand, if a remote or internal host accesses an FTP server on a host you are monitoring, the system can positively identify the application protocol.

If the system can identify the client used by a monitored host to connect to a non-monitored server, the system identifies the client's corresponding application protocol, but does not add the protocol to the network map. Note that client sessions must include a response from the server for application detection to occur.

The system characterizes each application that it detects; see Application Characteristics, on page 1736. The system uses these characteristics to create groups of applications, called *application filters*. Application filters are used to perform access control and to constrain search results and data used in reports and dashboard widgets.

You can also supplement application detector data using exported NetFlow records, Nmap active scans, and the host input feature.

Related Topics

Best Practices for Configuring Application Control, on page 1735
Application Detector Fundamentals, on page 2624

Application Detector Fundamentals

The system uses *application detectors* to identify the commonly used applications on your network. Use the Detectors page (**Policies** > **Application Detectors**) to view the detector list and customize detection capability.

Whether you can modify a detector or change its state (active or inactive) depends on its type. The system uses only active detectors to analyze application traffic.

Note Cisco-provided detectors may change with system and VDB updates. See the release notes and advisories for information on updated detectors.

Note For Firepower application identification, the ports are not listed intentionally. The application's associate ports are not reported for any of Cisco's applications because most of the applications are port-agnostic. Our platform's detection capabilities can identify services running at any port in the network.

Cisco-Provided Internal Detectors

Internal detectors are a special category of detectors for client, web application, and application protocol traffic. Internal detectors are delivered with system updates and are always on.

If an application matches against internal detectors designed to detect client-related activity and no specific client detector exists, a generic client may be reported.

Cisco-Provided Client Detectors

Client detectors detect client traffic and are delivered via VDB or system update, or are provided for import by Cisco Professional Services. You can activate and deactivate client detectors. You can export a client detector only if you import it.

Cisco-Provided Web Application Detectors

Web application detectors detect web applications in HTTP traffic payloads and are delivered via VDB or system update. Web application detectors are always on.

Cisco-Provided Application Protocol (Port) Detectors

Port-based application protocol detectors use well-known ports to identify network traffic. They are delivered via VDB or system update, or are provided for import by Cisco Professional Services. You can activate and deactivate application protocol detectors, and view a detector definition to use it as the basis for a custom detector.

Cisco-Provided Application Protocol (Firepower) Detectors

Firepower-based application protocol detectors analyze network traffic using Firepower application fingerprints and are delivered via VDB or system update. You can activate and deactivate application protocol detectors.

Custom Application Detectors

Custom application detectors are pattern-based. They detect patterns in packets from client, web application, or application protocol traffic. You have full control over imported and custom detectors.

Identification of Application Protocols in the Web Interface

The following table outlines how the system identifies detected application protocols:

Table 141: System Identification of Application Protocols

Identification	Description
application protocol name	The Firewall Management Center identifies an application protocol with its name if the application protocol was: • positively identified by the system • identified using NetFlow data and there is a port-application protocol correlation in `/etc/sf/services` • manually identified using the host input feature • identified by Nmap or another active source

Identification	Description
`pending`	The Firewall Management Center identifies an application protocol as `pending` if the system can neither positively nor negatively identify the application. Most often, the system needs to collect and analyze more connection data before it can identify a pending application. In the Application Details and Servers tables and in the host profile, the `pending` status appears only for application protocols where specific application protocol traffic was detected (rather than inferred from detected client or web application traffic).
`unknown`	The Firewall Management Center identifies an application protocol as `unknown` if: • the application does not match any of the system's detectors. • the application protocol was identified using NetFlow data, but there is no port-application protocol correlation in `/etc/sf/services`. • Snort has closed the session but it still persists in the device. Here, the traffic is allowed to pass through the firewall, but the application is not detected.
blank	All available detected data has been examined, but no application protocol was identified. In the Application Details and Servers tables and in the host profile, the application protocol is left blank for non-HTTP generic client traffic with no detected application protocol.

Implied Application Protocol Detection from Client Detection

If the system can identify the client used by a monitored host to access a non-monitored server, the Firewall Management Center infers that the connection is using the application protocol that corresponds with the client. (Because the system tracks applications only on monitored networks, connection logs usually do not include application protocol information for connections where a monitored host is accessing a non-monitored server.)

This process, or *implied application protocol detection*, has the following consequences:

- Because the system does not generate a New TCP Port or New UDP Port event for these servers, the server does not appear in the Servers table. In addition, you cannot trigger either discovery event alerts or correlation rules using the detection of these application protocol as a criterion.
- Because the application protocol is not associated with a host, you cannot view its details in host profiles, set its server identity, or use its information in host profile qualifications for traffic profiles or correlation rules. In addition, the system does not associate vulnerabilities with hosts based on this type of detection.

You can, however, trigger correlation events on whether the application protocol information is present in a connection. You can also use the application protocol information in connection logs to create connection trackers and traffic profiles.

Host Limits and Discovery Event Logging

When the system detects a client, server, or web application, it generates a discovery event unless the associated host has already reached its maximum number of clients, servers, or web applications.

Host profiles display up to 16 clients, 100 servers, and 100 web applications per host.

Note that actions dependent on the detection of clients, servers, or web applications are unaffected by this limit. For example, access control rules configured to trigger on a server will still log connection events.

Special Considerations for Application Detection

SFTP

In order to detect SFTP traffic, the same rule must also detect SSH.

Squid

The system positively identifies Squid server traffic when either:

- the system detects a connection from a host on your monitored network to a Squid server where proxy authentication is enabled, or
- the system detects a connection from a Squid proxy server on your monitored network to a target system (that is, the destination server where the client is requesting information or another resource).

However, the system cannot identify Squid service traffic if:

- a host on your monitored network connects to a Squid server where proxy authentication is disabled, or
- the Squid proxy server is configured to strip Via: header fields from its HTTP responses

SSL Application Detection

The system provides application detectors that can use session information from a Secure Socket Layers (SSL) session to identify the application protocol, client application, or web application in the session.

When the system detects an encrypted connection, it marks that connection as either a generic HTTPS connection or as a more specific secure protocol, such as SMTPS, when applicable. When the system detects an SSL session, it adds `SSL client` to the **Client** field in connection events for the session. If it identifies a web application for the session, the system generates discovery events for the traffic.

For SSL application traffic, managed devices can also detect the common name from the server certificate and match that against a client or web application from an SSL host pattern. When the system identifies a specific client, it replaces `SSL client` with the name of the client.

Because the SSL application traffic is encrypted, the system can use only information in the certificate for identification, not application data within the encrypted stream. For this reason, SSL host patterns can sometimes only identify the company that authored the application, so SSL applications produced by the same company may have the same identification.

In some instances, such as when an HTTPS session is launched from within an HTTP session, managed devices detect the server name from the client certificate in a client-side packet.

To enable SSL application identification, you must create access control rules that monitor responder traffic. Those rules must have either an application condition for the SSL application or URL conditions using the URL from the SSL certificate. For network discovery, the responder IP address does not have to be in the networks to monitor in the network discovery policy; the access control policy configuration determines whether the traffic is identified. To identify detections for SSL applications, you can filter by the `SSL protocol` tag, in the application detectors list or when adding application conditions in access control rules.

Referred Web Applications

Web servers sometimes refer traffic to other websites, which are often advertisement servers. To help you better understand the context for referred traffic occurring on your network, the system lists the web application that referred the traffic in the **Web Application** field in events for the referred session. The VDB contains a list of known referred sites. When the system detects traffic from one of those sites, the referring site is stored with the event for that traffic. For example, if an advertisement accessed via Facebook is actually hosted on Advertising.com, the detected Advertising.com traffic is associated with the Facebook web application. The system can also detect referring URLs in HTTP traffic, such as when a website provides a simple link to another site; in this case, the referring URL appears in the HTTP Referrer event field.

In events, if a referring application exists, it is listed as the web application for the traffic, while the URL is that for the referred site. In the example above, the web application for the connection event for that traffic would be Facebook, but the URL would be Advertising.com. A referred application may appear as the web application if no referring web application is detected, if the host refers to itself, or if there is a chain of referrals. In the dashboard, connection and byte counts for web applications include sessions where the web application is associated with traffic referred by that application.

Note that if you create a rule to act specifically on referred traffic, you should add a condition for the referred application, rather than the referring application. To block Advertising.com traffic referred from Facebook, for example, add an application condition to your access control rule for the Advertising.com application.

Application Detection in Snort 3

Note Snort 3 is now at parity with Snort 2, with respect to enabling AppID inspection exclusively on particular network subnets that are defined in the Network Discovery policy filters **if** no other configuration in the AC policy requires AppID to monitor all traffic.

In Snort 3, application detection is always enabled for all networks by default. To disable application detection, do the following:

Procedure

Step 1 Choose **Policies** > **Access Control heading** > **Access Control**, click edit policy and delete the application rules.

Step 2 Choose **Policies** > **Access Control heading** > **Decryption**, click delete to delete the SSL policy.

Step 3 Choose **Policies** > **Network Discovery**, click delete to delete the network discovery policy.

Step 4 Choose **Policies** > **Access Control heading** > **Access Control**, click **Edit** (✎) to the policy you want to edit and then click the **Security Intelligence** > **URLs** tab to delete the URLs Allow or Block list.

Step 5 As you cannot delete default DNS rules, choose **Policies** > **Access Control heading** > **DNS**, click edit and uncheck the enabled box to disable the DNS policy.

Step 6 In the access control policy, under the **Advanced** settings, disable the *Enable Threat Intelligence Director* and *Enable reputation enforcement on DNS traffic* options.

Step 7 Save and deploy the access control policy.

Requirements and Prerequisites for Application Detection

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Discovery Admin

Custom Application Detectors

If you use a custom application on your network, you can create a custom web application, client, or application protocol detector that provides the system with the information it needs to identify the application. The type of application detector is determined by your selections in the **Protocol**, **Type**, and **Direction** fields.

Client sessions must include a responder packet from the server for the system to begin detecting and identifying application protocols in server traffic. Note that, for UDP traffic, the system designates the source of the responder packet as the server.

If you have already created a detector on another Firewall Management Center, you can export it and then import it onto this Firewall Management Center. You can then edit the imported detector to suit your needs. You can export and import custom detectors as well as detectors provided by Cisco Professional Services. However, you **cannot** export or import any other type of Cisco-provided detectors.

Custom Application Detector and User-Defined Application Fields

You can use the following fields to configure custom application detectors and user-defined applications.

Custom Application Detector Fields: General

Use the following fields to configure basic and advanced custom application detectors.

Application Protocol

The application protocol you want to detect. This can be a system-provided application or a user-defined application.

If you want the application to be available for exemption from active authentication (configured in your identity rules), you must select or create an application protocol with the **`User-Agent Exclusion`** tag.

Description

A description for the application detector.

Name

A name for the application detector.

Detector Type

The type of detector, **Basic** or **Advanced**. Basic application detectors are created in the web interface as a series of fields. Advanced application detectors are created externally and uploaded as custom .lua files.

Custom Application Detector Fields: Detection Patterns

Use the following fields to configure the detection patterns for basic custom application detectors.

Direction

The source of the traffic the detector should inspect, **Client** or **Server**.

Offset

The location in a packet, in bytes from the beginning of the packet payload, where the system should begin searching for the pattern.

Because packet payloads start at byte 0, calculate the offset by subtracting 1 from the number of bytes you want to move forward from the beginning of the packet payload. For example, to look for the pattern in the fifth bit of the packet, type `4` in the **Offset** field.

Pattern

The pattern string associated with the **Type** you selected.

Ports

The port of the traffic the detector should inspect.

Protocol

The protocol you want to detect. Your protocol selection determines whether the **Type** or the **URL** field displays.

The protocol (and, in some cases, your subsequent selections in the **Type** and **Direction** fields) determine the type of application detector you create: web application, client, or application protocol.

Detector Type	Protocol	Type or Direction
Web Application	HTTP	**Type** is **Content Type** or **URL**
	RTMP	Any
	SSL	Any
Client	HTTP	**Type** is **User Agent**
	SIP	Any
	TCP or UDP	**Direction** is **Client**
Application Protocol	TCP or UDP	**Direction** is **Server**

Type

The type of pattern string you entered. The options you see are determined by the **Protocol** you selected. If you selected **RTMP** as the protocol, the **URL** field displays instead of the **Type** field.

Note If you select **User Agent** as the **Type**, the system automatically sets the **Tag** for the application to **User-Agent Exclusion**.

Type Selection	String Characteristics
Ascii	The string is ASCII encoded.
Common Name	The string is the value in the commonName field within the server response message.
Content Type	The string is the value in the content-type field within the server response header.
Hex	The string is in hexadecimal notation.
Organizational Unit	The string is the value in the organizationName field within the server response message.
SIP Server	The string is the value in the From field within the message header.
SSL Host	The string is the value in the server_name field within the ClientHello message.
URL	The string is a URL. **Note** The detector assumes that the string you enter is a complete section of the URL. For example, entering **cisco.com** would match **www.cisco.com/support** and **www.cisco.com**, but not **www.wearecisco.com**.
User Agent	The string is the value in the user-agent field within the GET request header. It is also available for the SIP protocol and indicates that the string is the value in the User-Agent field within the SIP message header.

URL

Either a full URL or a section of a URL from the swfURL field within the C2 message of a RTMP packet. This field displays instead of the **Type** field when you select **RTMP** as the **Protocol**.

Note The detector assumes that the string you enter is a complete section of the URL. For example, entering **cisco.com** would match **www.cisco.com/support** and **www.cisco.com**, but not **www.wearecisco.com**.

User-Defined Application Fields

Use the following fields to configure user-defined applications within basic and advanced custom application detectors.

Business Relevance

The likelihood that the application is used within the context of your organization's business operations, as opposed to recreationally: **Very High**, **High**, **Medium**, **Low**, or **Very Low**. Select the option that best describes the application.

Categories

A general classification for the application that describes its most essential function.

Description

A description for the application.

Name

A name for the application.

Risk

The likelihood that the application is used for purposes that might be against your organization's security policy: **Very High**, **High**, **Medium**, **Low**, or **Very Low**. Select the option that best describes the application.

Tags

One or more predefined tags that provide additional information about the application. If you want an application to be available for exemption from active authentication (configured in your identity rules), you must add the **User-Agent Exclusion** tag to your application.

Configuring Custom Application Detectors

You can configure basic or advanced custom application detectors.

Procedure

Step 1 Select **Policies** > **Application Detectors**.

Step 2 Click **Create Custom Detector**.

Step 3 Enter a **Name** and a **Description**.

Step 4 Choose an **Application Protocol** from the application drop-down list. You have the following options:

- If you are creating a detector for an existing application protocol (for example, if you want to detect a particular application protocol on a non-standard port), select the application protocol from the drop-down list.
- If you are creating a detector for a user-defined application, follow the procedure outlined in Create a User-Defined Application, on page 2633.

Step 5 Click **Detector Type** as **Basic** or **Advanced**.

Step 6 Click **OK**.

Step 7 Configure **Detection Patterns** or **Detection Criteria** or **Encrypted Visibility Engine Process Assignments**:

- If you are configuring a basic detector, specify preset **Detection Patterns** as described in Specifying Detection Patterns in Basic Detectors, on page 2634.
- If you are configuring an advanced detector, specify custom **Detection Criteria** as described in Specifying Detection Criteria in Advanced Detectors, on page 2635.
- If you are configuring an encrypted visibility engine (EVE) detector, specify custom EVE process assignments as described in *Specifying EVE Process Assignments* section in this chapter.

Caution
Advanced custom detectors are complex and require outside knowledge to construct valid .lua files. Incorrectly configured detectors could have a negative impact on performance or detection capability.

Step 8 Optionally, use **Packet Captures** to test the new detector as described in Testing a Custom Application Protocol Detector, on page 2637.

Step 9 Click **Save**.

Note
If you include the application in an access control rule, the detector is automatically activated and cannot be deactivated while in use.

What to do next

- Activate the detector as described in Activating and Deactivating Detectors, on page 2640.

Related Topics

Custom Application Detector and User-Defined Application Fields, on page 2629

Create a User-Defined Application

Applications, categories, and tags created here are available in access control rules and in the application filter object manager as well.

Caution Creating a user-defined application immediately restarts the Snort process without going through the deploy process. The system warns you that continuing restarts the Snort process and allows you to cancel; the restart occurs on any managed device in the current domain or in any of its child domains. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

Before you begin

- Begin configuring your custom application protocol detector as described in Configuring Custom Application Detectors, on page 2632.

Procedure

Step 1 On the **Create A Custom Application Detector** dialog box, click **Add** (+) next to the **Application** field.

Step 2 Type a **Name**.

Step 3 Type a **Description**.

Step 4 Select a **Business Relevance**.

Step 5 Select a **Risk**.

Step 6 Click **Add** next to Categories to add a category and type a new category name, or select an existing category from the **Categories** drop-down list.

Step 7 Click **Add** next to Tags to add a tag and type a new tag name, or select an existing tag from the **Tags** drop-down list.

Note
To associate the customized application detector domain for Policy Based Routing, you must select NSG from the **Tags** list.

Step 8 Click **OK**.

What to do next

- Continue configuring your custom application protocol detector as described in Configuring Custom Application Detectors, on page 2632. You must save and activate the detector before the system can use it to analyze traffic.

Related Topics

Custom Application Detector and User-Defined Application Fields, on page 2629

Specifying Detection Patterns in Basic Detectors

You can configure a custom application protocol detector to search application protocol packet headers for a particular pattern string. You can also configure detectors to search for multiple patterns; in that case the application protocol traffic must match all of the patterns for the detector to positively identify the application protocol.

Application protocol detectors can search for ASCII or hexadecimal patterns, using any offset.

Before you begin

- Begin configuring your custom application protocol detector as described in Configuring Custom Application Detectors, on page 2632.

Procedure

Step 1 On the **Create Detector** page, in the **Detection Patterns** section, click **Add**.

Step 2 Choose protocol type from the **Application** drop-down list.

Step 3 Choose pattern type from the **Type** drop-down list.

Step 4 Type a **Pattern** string that matches the **Type** you specified.

Step 5 Optionally, type the **Offset** (in bytes).

Step 6 Optionally, to identify application protocol traffic based on the port it uses, type a port from 1 to 65535 in the **Port(s)** field. To use multiple ports, separate them by commas.

Step 7 Click a **Direction**: `Client` or `Server`.

Step 8 Click **OK**.

Tip

If you want to delete a pattern, click **Delete** ( ) next to the pattern you want to delete.

What to do next

- Continue configuring your custom application protocol detector as described in Configuring Custom Application Detectors, on page 2632. You must save and activate the detector before the system can use it to analyze traffic.

Related Topics

Specifying Detection Criteria in Advanced Detectors, on page 2635

Specifying Detection Criteria in Advanced Detectors

Caution Advanced custom detectors are complex and require outside knowledge to construct valid .lua files. Incorrectly configured detectors could have a negative impact on performance or detection capability.

Caution Do not upload .lua files from untrusted sources.

Custom .lua files contain your custom application detector settings. Creating custom .lua files requires advanced knowledge of the lua programming language and experience with Cisco's C-lua API. Cisco strongly recommends you use the following to prepare .lua files:

- third-party instruction and reference material for the lua programming language
- The Open Source Detectors Developers Guide: https://www.snort.org/downloads
- OpenAppID Snort community resources: http://blog.snort.org/search/label/openappid

Note The system does not support .lua files that reference system calls or file I/O.

Before you begin

- Begin configuring your custom application protocol detector as described in Configuring Custom Application Detectors, on page 2632.
- Prepare to create a valid .lua file by downloading and studying the .lua files for comparable detectors. For more information about downloading detector files, see Viewing or Downloading Detector Details, on page 2638.
- Create a valid .lua file that contains your custom application detector settings.

Procedure

Step 1 On the **Create Detector** page for an advanced custom application detector, in the **Detection Criteria** section, click **Add**.

Step 2 Click **Browse...** to navigate to the `.lua` file and upload it.

Step 3 Click **OK**.

What to do next

- Continue configuring your custom application protocol detector as described in Configuring Custom Application Detectors, on page 2632. You must save and activate the detector before the system can use it to analyze traffic.

Related Topics

Specifying Detection Patterns in Basic Detectors, on page 2634

Specifying EVE Process Assignments

You can configure your own custom application detectors to map processes detected by the encrypted visibility engine (EVE) to new or existing applications.

Before you begin

- Begin configuring your custom application protocol detector as described in Configuring Custom Application Detectors, on page 2632.

Procedure

Step 1 On the Create Detector page, in the Encrypted Visibility Engine Process Assignments section, click **Add**.

Step 2 Enter the **Process Name** and **Minimum Process Confidence** value.

Note
You can enter text in the **Process Name** field and this is case-sensitive. The value should match the exact process name detected by EVE. The **Minimum Process Confidence** can be any number from 0 to 100. This is the number displayed in the **Encrypted Visibility Process Confidence Score** field in Connection Events.

For information about the **Encrypted Visibility Process Confidence Score** field, see the section *Connection and Security Intelligence Event Fields* in the Cisco Firepower Management Center Administration Guide.

Step 3 Click **Save**.

Step 4 In the Application Detector listing page, activate the detector that you created. For more information, see Activating and Deactivating Detectors, on page 2640. When you activate the detector, the detector files are pushed to all the FTDs registered on the Firewall Management Center.

What to do next

- Continue configuring your custom application protocol detector as described in Configuring Custom Application Detectors, on page 2632. You must save and activate the detector before the system can use it to analyze traffic.

Testing a Custom Application Protocol Detector

If you have a packet capture (pcap) file that contains packets with traffic from the application protocol you want to detect, you can test a custom application protocol detector against that pcap file. Cisco recommends using a simple, clean pcap file without unnecessary traffic.

Pcap files must be 256 KB or smaller; if you try to test your detector against a larger pcap file, the Firewall Management Center automatically truncates it and tests the incomplete file. You must fix the unresolved checksums in a pcap before using the file to test a detector.

Before you begin

- Configure your custom application protocol detector as described in Configuring Custom Application Detectors, on page 2632.

Procedure

Step 1 On the Create Detector page, in the Packet Captures section, click **Add**.

Step 2 Browse to the pcap file in the pop-up window and click **OK**.

Step 3 To test your detector against the contents of the pcap file, click evaluate next to the pcap file. A message indicates whether the test succeeded.

Step 4 Optionally, repeat steps 1 to 3 to test the detector against additional pcap files.

Tip

To delete a pcap file, click **Delete** ( ) next to the file you want to delete.

What to do next

- Continue configuring your custom application protocol detector as described in Configuring Custom Application Detectors, on page 2632. You must save and activate the detector before the system can use it to analyze traffic.

Viewing or Downloading Detector Details

You can use the detectors list to view application detector details (all detectors) and download detector details (custom application detectors only).

Procedure

Step 1 To view application detector details, do one of the following:

- See the *Cisco Firepower Application Detector Reference* for the relevant VDB version at https://www.cisco.com/c/en/us/support/security/defense-center/products-technical-reference-list.html
- a. Select **Policies** > **Application Detectors**.

  b. Filter the list to find a particular detector.

  c. Click **Information**( )

Step 2 To download detector details for a custom application detector, click **Download** ( ).

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have the necessary permissions.

Sorting the Detector List

By default, the Detectors page lists detectors alphabetically by name. An up or down arrow next to a column heading indicates that the page is sorted by that column in that direction.

Procedure

Step 1 Select **Policies** > **Application Detectors**.

Step 2 Click the appropriate column heading.

Filtering the Detector List

Procedure

Step 1 Select **Policies** > **Application Detectors**.

Step 2 Expand one of the filter groups described in Filter Groups for the Detector List, on page 2639 and select the check box next to a filter. To select all filters in a group, right-click the group name and select **Check All**.

Step 3 If you want to remove a filter, click **Remove** (✕) in the name of the filter in the **Filters** field or disable the filter in the filter list. To remove all filters in a group, right-click the group name and select **Uncheck All**.

Step 4 If you want to remove all filters, click **Clear all** next to the list of filters applied to the detectors.

Filter Groups for the Detector List

You can use several filter groups, separately or in combination, to filter the list of detectors.

Name

Finds detectors with names or descriptions containing the string you type. Strings can contain any alphanumeric or special character.

Custom Filter

Finds detectors matching a custom application filter created on the object management page.

Author

Finds detectors according to who created the detector. You can filter detectors by:

- any individual user who has created or imported a custom detector
- Cisco, which represents all Cisco-provided detectors *except* individually imported add-on detectors (you are the author for any detector that you import)
- **Any User**, which represents all detectors not provided by Cisco

State

Finds detectors according to their state, that is, **Active** or **Inactive**.

Type

Finds detectors according to the detector type, as described in Application Detector Fundamentals, on page 2624.

Protocol

Finds detectors according to which traffic protocol the detector inspects.

Category

Finds detectors according to the categories assigned to the application they detect.

Tag

Finds detectors according to the tags assigned to the application they detect.

Risk

Finds detectors according to the risks assigned to the application they detect: **Very High**, **High**, **Medium**, **Low**, and **Very Low**.

Business Relevance

Finds detectors according to the business relevance assigned to the application they detect: **Very High**, **High**, **Medium**, **Low**, and **Very Low**.

Navigating to Other Detector Pages

Procedure

Step 1 Select **Policies** > **Application Detectors**.

Step 2 If you want to view the next page, click **Right Arrow** ( ).

Step 3 If you want to view the previous page, click **Left Arrow** ( ).

Step 4 If you want to view a different page, type the page number and press Enter.

Step 5 If you want to jump to the last page, click **Right End Arrow** ( ).

Step 6 If you want to jump to the first page, click **Left End Arrow** ( ).

Activating and Deactivating Detectors

You must activate a detector before you can use it to analyze network traffic. By default, all Cisco-provided detectors are activated.

You can activate multiple application detectors for each port to supplement the system's detection capability.

When you include an application in an access control rule in a policy and that policy is deployed, if there is no active detector for that application, one or more detectors automatically activate. Similarly, while an application is in use in a deployed policy, you cannot deactivate a detector if deactivating leaves no active detectors for that application.

Tip For improved performance, deactivate any application protocol, client, or web application detectors you do not intend to use.

Caution Activating or deactivating a system or custom application detector immediately restarts the Snort process without going through the deploy process. The system warns you that continuing restarts the Snort process and allows you to cancel; the restart occurs on any managed device in the current domain or in any of its child domains. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

Procedure

Step 1 Select **Policies** > **Application Detectors**.

Step 2 Click the slider next to the detector you want to activate or deactivate. If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Note
Some application detectors are required by other detectors. If you deactivate one of these detectors, a warning appears to indicate that the detectors that depend on it are also disabled.

Editing Custom Application Detectors

Use the following procedure to modify custom application detectors.

Procedure

Step 1 Select **Policies** > **Application Detectors**.

Step 2 Click **Edit** (✎) next to the detector you want to modify. If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Make changes to the detector as described in Configuring Custom Application Detectors, on page 2632.

Step 4 You have the following saving options, depending on the state of the detector:

- To save an inactive detector, click **Save**.
- To save an inactive detector as a new, inactive detector, click **Save as New**.
- To save an active detector and immediately start using it, click **Save and Reactivate**.

Caution
Saving and reactivating a custom application detector immediately restarts the Snort process without going through the deploy process. The system warns you that continuing restarts the Snort process and allows you to cancel; the restart occurs on any managed device in the current domain or in any of its child domains. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

- To save an active detector as a new, inactive detector, click **Save as New**.

Deleting Detectors

You can delete custom detectors as well as individually imported add-on detectors provided by Cisco Professional Services. You cannot delete any of the other Cisco-provided detectors, though you can deactivate many of them.

Note While a detector is in use in a deployed policy, you cannot delete the detector.

Caution Deleting an activated custom application detector immediately restarts the Snort process without going through the deploy process. The system warns you that continuing restarts the Snort process and allows you to cancel; the restart occurs on any managed device in the current domain or in any of its child domains. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

Procedure

Step 1 Select **Policies** > **Application Detectors**.

Step 2 Click **Delete** ( ) next to the detector you want to delete. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 3 Click **OK**.

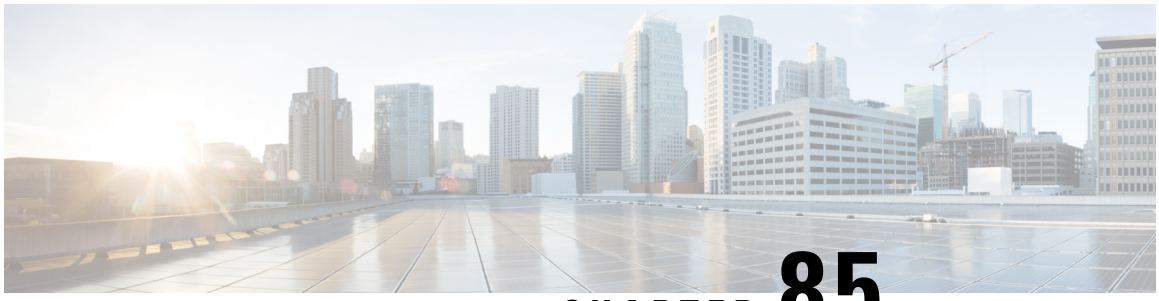

CHAPTER 85

Network Discovery Policies

The following topics describe how to create, configure, and manage network discovery policies:

- Overview: Network Discovery Policies, on page 2643
- Requirements and Prerequisites for Network Discovery Policies, on page 2644
- Network Discovery Customization, on page 2644
- Network Discovery Rules, on page 2645
- Configuring Advanced Network Discovery Options, on page 2654
- Troubleshooting Your Network Discovery Strategy, on page 2662

Overview: Network Discovery Policies

The network discovery policy on the Firewall Management Center controls how the system collects data on your organization's network assets and which network segments and ports are monitored.

Discovery rules within the policy specify which networks and ports the system monitors to generate discovery data based on network data in traffic, and the zones to which the policy is deployed. Within a rule, you can configure whether hosts, applications, and non-authoritative users are discovered. You can create rules to exclude networks and zones from discovery. You can configure discovery of data from NetFlow exporters and restrict the protocols for traffic where user data is discovered on your network.

The network discovery policy has a single default rule in place, configured to discover applications from all observed traffic. The rule does not exclude any networks, zones, or ports, host and user discovery is not configured, and the rule is not configured to monitor a NetFlow exporter. This policy is deployed by default to any managed devices when they are registered to the Firewall Management Center. To begin collecting host or user data, you must add or modify discovery rules and re-deploy the policy to a device.

If you want to adjust the scope of network discovery, you can create additional discovery rules and modify or remove the default rule.

Remember that the access control policy for each managed device defines the traffic that you permit for that device and, therefore, the traffic you can monitor with network discovery. If you block certain traffic using access control, the system cannot examine that traffic for host, user, or application activity. For example, if an access control policy blocks access to social networking applications, the system cannot provide any discovery data on those applications.

If you enable traffic-based user detection in your discovery rules, you can detect non-authoritative users through user login activity in traffic over a set of application protocols. You can disable discovery in particular protocols across all rules if needed. Disabling some protocols can help avoid reaching the user limit associated with your Firewall Management Center model, reserving available user count for users from the other protocols.

Advanced network discovery settings allow you to manage what data is logged, how discovery data is stored, what indications of compromise (IOC) rules are active, what vulnerability mappings are used for impact assessment, and what happens when sources offer conflicting discovery data. You can also add sources for host input and NetFlow exporters to monitor.

Requirements and Prerequisites for Network Discovery Policies

Model Support

Any.

Supported Domains

Leaf

User Roles

- Admin
- Discovery Admin

Network Discovery Customization

The information about your network traffic collected by the system is most valuable to you when the system can correlate this information to identify the hosts on your network that are most vulnerable and most important.

As an example, if you have several devices on your network running a customized version of SuSE Linux, the system cannot identify that operating system and so cannot map vulnerabilities to the hosts. However, knowing that the system has a list of vulnerabilities for SuSE Linux, you may want to create a custom fingerprint for one of the hosts that can then be used to identify the other hosts running the same operating system. You can include a mapping of the vulnerability list for SuSE Linux in the fingerprint to associate that list with each host that matches the fingerprint.

The system also allows you to input host data from third-party systems directly into the network map, using the host input feature. However, third-party operating system or application data does not automatically map to vulnerability information. If you want to see vulnerabilities and perform impact correlation for hosts using third-party operating system, server, and application protocol data, you must map the vendor and version information from the third-party system to the vendor and version listed in the vulnerability database (VDB). You also may want to maintain the host input data on an ongoing basis. Note that even if you map application data to system vendor and version definitions, imported third-party vulnerabilities are not used for impact assessment for clients or web applications.

If the system cannot identify application protocols running on hosts on your network, you can create user-defined application protocol detectors that allow the system to identify the applications based on a port or a pattern. You can also import, activate, and deactivate certain application detectors to further customize the application detection capability.

You can also replace detection of operating system and application data using scan results from the Nmap active scanner or augment the vulnerability lists with third-party vulnerabilities. The system may reconcile data from multiple sources to determine the identity for an application.

Configuring the Network Discovery Policy

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Configure the following components of your policy:

- Discovery rules — See Configuring Network Discovery Rules, on page 2646.
- Traffic-based detection for users — See Configuring Traffic-Based User Detection, on page 2653.
- Advanced network discovery options — See Configuring Advanced Network Discovery Options, on page 2654.
- Custom operating system definitions (fingerprints) — See Creating a Custom Fingerprint for Clients, on page 2592 and Creating a Custom Fingerprint for Servers, on page 2594.

Network Discovery Rules

Network discovery rules allow you to tailor the information discovered for your network map to include only the specific data you want. Rules in your network discovery policy are evaluated sequentially. You can create rules with overlapping monitoring criteria, but doing so may affect your system performance.

When you exclude a host or a network from monitoring, the host or network does not appear in the network map and no events are reported for it. However, when the host discovery rules for the local IP are disabled, the detection engine instances are impacted by a higher processing load, as it builds data from each flow afresh rather than using the existing host data.

We recommend that you exclude load balancers (or specific ports on load balancers) and NAT devices from monitoring. These devices may create excessive and misleading events, filling the database and overloading the Firewall Management Center. For example, a monitored NAT device might exhibit multiple updates of its operating system in a short period of time. If you know the IP addresses of your load balancers and NAT devices, you can exclude them from monitoring.

Tip The system can identify many load balancers and NAT devices by examining your network traffic.

In addition, if you need to create a custom server fingerprint, you should temporarily exclude from monitoring the IP address that you are using to communicate with the host you are fingerprinting. Otherwise, the network map and discovery event views will be cluttered with inaccurate information about the host represented by that IP address. After you create the fingerprint, you can configure your policy to monitor that IP address again.

Cisco also recommends that you **not** monitor the same network segment with NetFlow exporters and managed devices. Although ideally you should configure your network discovery policy with non-overlapping rules, the system does drop duplicate connection logs generated by managed devices. However, you **cannot** drop duplicate connection logs for connections detected by both a managed device and a NetFlow exporter.

Configuring Network Discovery Rules

You can configure discovery rules to tailor the discovery of host and application data to your needs.

Tip In most cases, we recommend restricting discovery to the addresses in RFC 1918.

Before you begin

- Make sure you are logging connections for the traffic where you want to discover network data ; see *Best Practices for Connection Logging* in the Cisco Secure Firewall Management Center Administration Guide.
- If you want to collect exported NetFlow records, add a NetFlow Exporter as described in Adding NetFlow Exporters to a Network Discovery Policy, on page 2658.
- If you will want to view discovery performance graphs, you must enable hosts, users, and applications in your discovery rule. Note that this may impact system performance.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Add Rule**.

Step 3 Set the **Action** for the rule as described in Actions and Discovered Assets, on page 2646.

Step 4 Set optional discovery parameters:

- Restrict the rule action to specific networks; see Restricting the Monitored Network, on page 2648.
- Restrict the rule action to traffic in specific zones; see Configuring Zones in Network Discovery Rules, on page 2651.
- Exclude ports from monitoring; see Excluding Ports in Network Discovery Rules, on page 2650.
- Configure the rule for NetFlow data discovery; see Configuring Rules for NetFlow Data Discovery, on page 2648.

Step 5 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Actions and Discovered Assets

When you configure a discovery rule, you must select an action for the rule. The effect of that action depends on whether you are using the rule to discover data from a managed device or from a NetFlow exporter.

The following table describes what assets are discovered by rules with the specified action settings in those two scenarios.

Table 142: Discovery Rule Actions

Action	Option	Managed Device	NetFlow Exporter
Exclude	--	Excludes the specified network from monitoring. If the source or destination host for a connection is excluded from discovery, the connection is recorded but discovery events are not created for excluded hosts.	Excludes the specified network from monitoring. If the source or destination host for a connection is excluded from discovery, the connection is recorded but discovery events are not created for excluded hosts.
Discover	Hosts	Adds hosts to the network map based on discovery events. (Optional, unless user discovery is enabled, then required.)	Adds hosts to the network map and logs connections based on NetFlow records. (Required)
Discover	Applications	Adds applications to the network map based on application detectors. Note that you cannot discover hosts or users in a rule without also discovering applications. (Required)	Adds application protocols to the network map based on NetFlow records and the port-application protocol correlation in `/etc/sf/services`. (Optional)
Discover	Users	Adds users to the users table and logs user activity based on traffic-based detection on the user protocols configured in the network discovery policy. (Optional)	n/a
Log NetFlow Connections	--	n/a	Logs NetFlow connections only. Does not discover hosts or applications.

If you want the rule to monitor managed device traffic, application logging is required. If you want the rule to monitor users, host logging is required. If you want the rule to monitor exported NetFlow records, you cannot configure it to log users, and logging applications is optional.

Note The system detects connections in exported NetFlow records based on the **Action** settings in the network discovery policy. The system detects connections in managed device traffic based on access control policy settings.

Monitored Networks

A discovery rule causes discovery of monitored assets only in traffic to and from hosts in the specified networks. For a discovery rule, discovery occurs for connections that have at least one IP address within the networks specified, with events generated only for IP addresses within the networks to monitor. The default discovery rule discovers applications from all observed traffic (`0.0.0.0/0` for all IPv4 traffic, and `::/0` for all IPv6 traffic).

If you configure a rule to handle NetFlow discovery and log only connections data, the system also logs connections to and from IP addresses in the specified networks. Note that network discovery rules provide the only way to log NetFlow network connections.

You can also use network object or object groups to specify the networks to monitor.

Restricting the Monitored Network

Every discovery rule must include at least one network.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Add Rule**.

Step 3 Click **Networks**, if it is not already open.

Step 4 (Optional) Add network objects to the Available Networks list as described in Creating Network Objects During Discovery Rule Configuration, on page 2649.

If you modify a network object used in the network discovery policy, the changes do not take effect for discovery until you deploy the configuration changes.

Step 5 Specify a network:

- Choose a network from the **Available Networks** list. If the network does not immediately appear on the list, click **Reload** ( ).
- Enter the IP address into the text box below the Available Networks label.

Step 6 Click **Add**.

Step 7 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Configuring Rules for NetFlow Data Discovery

The system can use data from NetFlow exporters to generate connection and discovery events, and to add host and application data to the network map.

If you choose a NetFlow exporter in a discovery rule, the rule is limited to discovery of NetFlow data for the specified networks. Choose the NetFlow device to monitor before you configure other aspects of rule behavior, as the available rule actions change when you choose a NetFlow device. You cannot configure port exclusions for monitoring NetFlow exporters.

Before you begin

- Add NetFlow-enabled devices to the network discovery policy; see Adding NetFlow Exporters to a Network Discovery Policy, on page 2658.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Add Rule**.

Step 3 Choose **NetFlow Device**.

Step 4 From the **NetFlow Device** drop-down list, choose the IP address of the NetFlow exporter to be monitored.

Step 5 Specify the type of NetFlow data you want the system managed device to collect:

- Connection only — Choose `Log NetFlow Connections` from the **Action** drop-down list.
- Host, Application, and Connection — Choose `Discover` from the **Action** drop-down list. The system automatically checks the **Hosts** check box and enables collection of connection data. Optionally, you can check the **Application** check box to collect application data.

Step 6 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Creating Network Objects During Discovery Rule Configuration

You can add new network objects to the list of available networks that appears in a discovery rule by adding them to the list of reusable network objects and groups.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 In **Networks**, click **Add Rule**.

Step 3 Click **Add** (+) next to **Available Networks**.

Step 4 Create a network object as described in Creating Network Objects, on page 1403.

Step 5 Finish adding the network discovery rule as described in Configuring Network Discovery Rules, on page 2646.

Port Exclusions

Just as you can exclude hosts from monitoring, you can exclude specific ports from monitoring. For example:

- Load balancers can report multiple applications on the same port in a short period of time. You can configure your network discovery rules so that they exclude that port from monitoring, such as excluding port 80 on a load balancer that handles a web farm.
- Your organization may use a custom client that uses a specific range of ports. If the traffic from this client generates excessive and misleading events, you can exclude those ports from monitoring. Similarly, you may decide that you do not want to monitor DNS traffic. In that case, you could configure your rules so that your discovery policy does not monitor port 53.

When adding ports to exclude, you can decide whether to use a reusable port object from the Available Ports list, add ports directly to the source or destination exclusion lists, or create a new reusable port and then move it into the exclusion lists.

Note You cannot exclude ports in rules handling NetFlow data discovery.

Excluding Ports in Network Discovery Rules

You cannot exclude ports in rules handling NetFlow data discovery.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Add Rule**.

Step 3 Click **Port Exclusions**.

Step 4 Optionally, add port objects to the Available Ports list as described in Creating Port Objects During Discovery Rule Configuration, on page 2650.

Step 5 Exclude specific source ports from monitoring, using either of the following methods:

- Choose a port or ports from the **Available Ports** list and click **Add to Source**.
- To exclude traffic from a specific source port without adding a port object, under the **Selected Source Ports** list, choose a **Protocol**, enter a **Port** number (a value from 1 to 65535), and click **Add**.

Step 6 Exclude specific destination ports from monitoring, using either of the following methods:

- Choose a port or ports from the **Available Ports** list and click **Add to Destination**.
- To exclude traffic from a specific destination port without adding a port object, under the **Selected Destination Ports** list, choose a **Protocol**, enter a **Port** number, and click **Add**.

Step 7 Click **Save** to save the changes you made.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Creating Port Objects During Discovery Rule Configuration

You can add new port objects to the list of available ports that appears in a discovery rule by adding them to the list of reusable port objects and groups that can be used anywhere in the system.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 In Networks, click **Add Rule**.

Step 3 Click **Port Exclusions**.

Step 4 To add a port to the Available Ports list, click **Add** (+).

Step 5 Enter **Name**.

Step 6 In the **Protocol** field, specify the protocol of the traffic you want to exclude.

Step 7 In the **Port** field, enter the ports you want to exclude from monitoring.

You can specify a single port, a range of ports using the dash (-), or a comma-separated list of ports and port ranges. Allowed port values are from 1 to 65535.

Step 8 Click **Save**.

Step 9 If the port does not immediately appear on the list, click **Refresh**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Zones in Network Discovery Rules

To improve performance, discovery rules can be configured so that the zones in the rule include the sensing interfaces on your managed devices that are physically connected to the networks-to-monitor in the rule.

Unfortunately, you may not always be kept informed of network configuration changes. A network administrator may modify a network configuration through routing or host changes without informing you, which may make it challenging to stay on top of proper network discovery policy configurations. If you do not know how the sensing interfaces on your managed devices are physically connected to your network, leave the zone configuration as the default. This default causes the system to deploy the discovery rule to all zones in your deployment. (If no zones are excluded, the system deploy the discovery policy to all zones.)

Configuring Zones in Network Discovery Rules

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Add Rule**.

Step 3 Click **Zones**.

Step 4 Choose a zone or zones from the **Available Zones** list.

Step 5 Click **Save** to save the changes you made.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

The Traffic-Based Detection Identity Source

Traffic-based detection is the only non-authoritative identity source supported by the system. When configured, managed devices detect LDAP, AIM, POP3, IMAP, Oracle, SIP (VoIP), FTP, HTTP, MDNS, and SMTP logins on the networks you specify. The data gained from traffic-based detection can be used only for user awareness. Unlike authoritative identity sources, you configure traffic-based detection in your network discovery policy as described in Configuring Traffic-Based User Detection, on page 2653.

Note the following limitations:

- Traffic-based detection interprets only Kerberos logins for LDAP connections as LDAP authentications. Managed devices cannot detect encrypted LDAP authentications using protocols such as SSL or TLS.
- Traffic-based detection detects AIM logins using the OSCAR protocol only. They cannot detect AIM logins using TOC2.
- Traffic-based detection cannot restrict SMTP logging. This is because users are not added to the database based on SMTP logins; although the system detects SMTP logins, the logins are not recorded unless there is already a user with a matching email address in the database.

Traffic-based detection also records failed login attempts. A failed login attempt does not add a new user to the list of users in the database. The user activity type for detected failed login activity detected by traffic-based detection is **Failed User Login**.

Note The system cannot distinguish between failed and successful HTTP logins. To see HTTP user information, you must enable **Capture Failed Login Attempts** in the traffic-based detection configuration.

Caution Enabling or disabling non-authoritative, traffic-based user detection over the HTTP, FTP, or MDNS protocols, using the network discovery policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

Traffic-Based Detection Data

When a device detects a login using traffic-based detection, it sends the following information to the Firewall Management Center to be logged as user activity:

- The user name identified in the login.
- The time of the login.
- The IP address involved in the login, which can be the IP address of the user's host (for LDAP, POP3, IMAP, and AIM logins), the server (for HTTP, MDNS, FTP, SMTP and Oracle logins), or the session originator (for SIP logins).
- The user's email address (for POP3, IMAP, and SMTP logins).
- The name of the device that detected the login.

If the user was previously detected, the Firewall Management Center updates that user's login history. Note that the Firewall Management Center can use the email addresses in POP3 and IMAP logins to correlate with LDAP users. This means that, for example, if the Firewall Management Center detects a new IMAP login, and the email address in the IMAP login matches that for an existing LDAP user, the IMAP login does not create a new user; rather, it updates the LDAP user's history.

If the user was previously undetected, the Firewall Management Center adds the user to the users database. Unique AIM, SIP, and Oracle logins always create new user records, because there is no data in those login events that the Firewall Management Center can correlate with other login types.

The Firewall Management Center does **not** log user activity or user identities in the following cases:

- If you configured the network discovery policy to ignore that login type
- If a managed device detects an SMTP login, but the users database does not contain a previously detected LDAP, POP3, or IMAP user with a matching email address

The user data is added to the users table.

Traffic-Based Detection Strategies

You can restrict the protocols where user activity is discovered to reduce the total number of detected users so you can focus on users likely to provide the most complete user information. Limiting protocol detection helps minimize user name clutter and preserve storage space on your Firewall Management Center.

Consider the following when selecting traffic-based detection protocols:

- Obtaining user names through protocols such as AIM, POP3, and IMAP may introduce user names not relevant to your organization due to network access from contractors, visitors, and other guests.
- AIM, Oracle, and SIP logins may create extraneous user records. This occurs because these login types are not associated with any of the user metadata that the system obtains from an LDAP server, nor are they associated with any of the information contained in the other types of login that your managed devices detect. Therefore, the Firewall Management Center cannot correlate these users with other types of users.

Configuring Traffic-Based User Detection

When you enable traffic-based user detection in a network discovery rule, host discovery is automatically enabled. For more information about traffic-based detection, see The Traffic-Based Detection Identity Source, on page 2651.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Users**.

Step 3 Click **Edit** ( ).

Step 4 Check the check boxes for protocols where you want to detect logins or clear check boxes for protocols where you do not want to detect logins, and choose whether you want to **Capture Failed Login Attempts**.

Step 5 Click **Save**.

What to do next

Caution Enabling or disabling non-authoritative, traffic-based user detection over the HTTP, FTP, or MDNS protocols, using the network discovery policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior, on page 254 for more information.

- Configure network discovery rules to discover users as described in Configuring Network Discovery Rules, on page 2646.
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Configuring Advanced Network Discovery Options

The Advanced of the network discovery policy allows you to configure policy-wide settings for what events are detected, how long discovery data is retained and how often it is updated, what vulnerability mappings are used for impact correlation, and how operating system and server identity conflicts are resolved. In addition, you can add host input sources and NetFlow exporters to allow import of data from other sources.

Note Database event limits for discovery and user activity events are set in system configuration.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Edit** (✎) or **Add** (+) next to the setting you want to modify:

- Data Storage Settings — Update the settings as described in Configuring Network Discovery Data Storage, on page 2660.
- Event Logging Settings — Update the settings as described in Configuring Network Discovery Event Logging, on page 2661.
- General Settings — Update the settings as described in Configuring Network Discovery General Settings, on page 2655.
- Identity Conflict Settings — Update the settings as described in Configuring Network Discovery Identity Conflict Resolution, on page 2656.
- Indications of Compromise Settings — Update the settings as described in Enabling Indications of Compromise Rules, on page 2658.
- NetFlow Exporters — Update the settings as described in Adding NetFlow Exporters to a Network Discovery Policy, on page 2658.
- OS and Server Identity Sources — Update the settings as described in Adding Network Discovery OS and Server Identity Sources, on page 2661.
- Vulnerabilities to use for Impact Assessment — Update the settings as described in Enabling Network Discovery Vulnerability Impact Assessment, on page 2657.

Step 4 Click **Save**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Network Discovery General Settings

The general settings control how often the system updates network maps and whether server banners are captured during discovery.

Capture Banners

Select this check box if you want the system to store header information from network traffic that advertises server vendors and versions (“banners”). This information can provide additional context to the information gathered. You can access server banners collected for hosts by accessing server details.

Update Interval

The interval at which the system updates information (such as when any of a host’s IP addresses was last seen, when an application was used, or the number of hits for an application). The default setting is 3600 seconds (1 hour).

Note that setting a lower interval for update timeouts provides more accurate information in the host display, but generates more network events.

Configuring Network Discovery General Settings

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Edit** (✎) next to **General Settings**.

Step 4 Update the settings as described in Network Discovery General Settings, on page 2655.

Step 5 Click **Save** to save the general settings.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Network Discovery Identity Conflict Settings

The system determines which operating system and applications are running on a host by matching fingerprints for operating systems and servers against patterns in traffic. To provide the most reliable operating system and server identity information, the system collates fingerprint information from several sources.

The system uses all passive data to derive operating system identities and assign a confidence value.

By default, unless there is an identity conflict, identity data added by a scanner or third-party application overrides identity data detected by the system. You can use the Identity Sources settings to rank scanner and third-party application fingerprint sources by priority. The system retains one identity for each source, but only data from the highest priority third-party application or scanner source is used as the current identity. Note, however, that user input data overrides scanner and third-party application data regardless of priority.

An identity conflict occurs when the system detects an identity that conflicts with an existing identity that came from either the active scanner or third-party application sources listed in the Identity Sources settings or from a system user. By default, identity conflicts are not automatically resolved and you must resolve them through the host profile or by rescanning the host or re-adding new identity data to override the passive identity. However, you can set your system to automatically resolve the conflict by keeping either the passive identity or the active identity.

Generate Identity Conflict Event

Specifies whether the system generates an event when an identity conflict occurs.

Automatically Resolve Conflicts

From the **Automatically Resolve Conflicts** drop-down list, choose one of the following:

- **Disabled** if you want to force manual conflict resolution of identity conflicts
- **Identity** if you want to the system to use the passive fingerprint when an identity conflict occurs
- **Keep Active** if you want the system to use the current identity from the highest priority active source when an identity conflict occurs

Configuring Network Discovery Identity Conflict Resolution

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Edit** (✎) next to **Identity Conflict Settings**.

Step 4 Update the settings in the Edit Identity Conflict Settings pop-up window as described in Network Discovery Identity Conflict Settings, on page 2655.

Step 5 Click **Save** to save the identity conflict settings.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Network Discovery Vulnerability Impact Assessment Options

You can configure how the system performs impact correlation with intrusion events. Your choices are as follows:

- Check the **Use Network Discovery Vulnerability Mappings** check box if you want to use system-based vulnerability information to perform impact correlation.
- Check the **Use Third-Party Vulnerability Mappings** check box if you want to use third-party vulnerability references to perform impact correlation. For more information, see the *Firepower System Host Input API Guide*.

You can check either or both of the check boxes. If the system generates an intrusion event and the host involved in the event has servers or an operating system with vulnerabilities in the selected vulnerability mapping sets, the intrusion event is marked with the Vulnerable (level 1: red) impact icon. For any servers which do not have vendor or version information, note that you need to enable vulnerability mapping in the Firewall Management Center configuration.

If you clear both check boxes, intrusion events will **never** be marked with the Vulnerable (level 1: red) impact icon.

Related Topics

Mapping Third-Party Vulnerabilities, on page 2600

Enabling Network Discovery Vulnerability Impact Assessment

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Edit** (✎) next to **Vulnerabilities to use for Impact Assessment**.

Step 4 Update the settings in the Edit Vulnerability Settings pop-up window as described in Network Discovery Vulnerability Impact Assessment Options, on page 2656.

Step 5 Click **Save** to save the vulnerability settings.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Indications of Compromise

The system uses IOC rules in the network discovery policy to identify a host as likely to be compromised by malicious means. When a host meets the conditions specified in these system-provided rules, the system tags it with an *indication of compromise* (IOC). The related rules are known as *IOC rules*. Each IOC rule corresponds to one type of IOC tag. The *IOC tags* specify the nature of the likely compromise.

The Firewall Management Center can tag the host and user involved when one of the following things occurs:

- The system correlates data gathered about your monitored network and its traffic, using intrusion, connection, Security Intelligence, and file or malware events, and determines that a potential IOC has occurred.
- The Firewall Management Center can import IOC data from your Secure Endpoint deployments via the AMP cloud. Because this data examines activity on a host itself—such as actions taken by or on individual programs—it can provide insights into possible threats that network-only data cannot. For your convenience, the Firewall Management Center automatically obtains any new IOC tags that Cisco develops from the AMP cloud.

To configure this feature, see Enabling Indications of Compromise Rules, on page 2658.

You can also write correlation rules against host IOC data and compliance allow lists that account for IOC-tagged hosts.

To investigate and work with tagged IOCs, see Cisco Secure Firewall Management Center Administration Guide.

Enabling Indications of Compromise Rules

For your system to detect and tag indications of compromise (IOC), you must first activate at least one IOC rule in your network discovery policy. Each IOC rule corresponds to one type of IOC tag, and all IOC rules are predefined by Cisco; you cannot create original rules. You can enable any or all rules, depending on the needs of your network and organization. For example, if hosts using software such as Microsoft Excel never appear on your monitored network, you may decide not to enable the IOC tags that pertain to Excel-based threats.

Tip To disable IOC rules for individual hosts or their associated users, see the *Discovery Events* chapter in the Cisco Secure Firewall Management Center Administration Guide.

Before you begin

Because IOC rules trigger based on data provided by other components of the system and by Secure Endpoint, those components must be correctly licensed and configured for IOC rules to set IOC tags. Enable the system features associated with the IOC rules you will enable, such as intrusion detection and prevention (IPS) and Advanced Malware Protection (AMP). If an IOC rule's associated feature is not enabled, no relevant data is collected and the rule cannot trigger.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Edit** (✎) next to **Indications of Compromise Settings**.

Step 4 To toggle the entire IOC feature off or on, click the slider next to **Enable IOC**.

Step 5 To globally enable or disable individual IOC rules, click the slider in the rule's **Enabled** column.

Step 6 Click **Save** to save your IOC rule settings.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Adding NetFlow Exporters to a Network Discovery Policy

Use this procedure to add NetFlow exporters. Note that you cannot delete an exporter if it is currently in use in a disocvery rule.

Before you begin

- Review prerequisites: Requirements for Using NetFlow Data, on page 2584
- Configure NetFlow exporters: NetFlow Data, on page 2583

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Add** (+) next to **NetFlow Devices**.

Step 4 Enter the **IP Address** of the exporter.

Step 5 Click **Save**.

What to do next

- Configure a network discovery rule to monitor NetFlow traffic: Configuring Network Discovery Rules, on page 2646
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Network Discovery Data Storage Settings

Discovery data storage settings include the host limit and timeout settings.

When Host Limit Reached

The number of hosts a Secure Firewall Management Center can monitor, and therefore store in the network map, depends on its model. The **When Host Limit Reached** option controls what happens when you detect a new host after you reach the host limit. You can:

Drop hosts

The system drops the host that has remained inactive for the longest time, then adds the new host. This is the default setting.

Don't insert new hosts

The system does not track any newly discovered hosts. The system only tracks new hosts after the host count drops below the limit, such as after an administrator increases the domain's host limit or manually deletes hosts from the network map, or if the system identifies hosts as timed-out due to inactivity.

Table 143: Reaching the Host Limit with Multitenancy

Setting	Domain Host Limit Set?	Domain Host Limit Reached	Ancestor Domain Host Limit Reached
Drop hosts	yes	Drops oldest host in the constrained domain.	Drops the oldest host among all descendant leaf domains configured to drop hosts. If no host can be dropped, does not add the host.
	no	n/a	Drops the oldest host among all descendant leaf domains configured to drop hosts and that share the general pool.
Don't insert new hosts	yes or no	Does not add the host.	Does not add the host.

Host Timeout

The amount of time that passes, in minutes, before the system drops a host from the network map due to inactivity. The default setting is 10080 minutes (one week). Individual host IP and MAC addresses can time out individually, but a host does not disappear from the network map unless all its associated addresses time out.

To avoid premature timeout of hosts, make sure that the host timeout value is longer than the update interval in the network discovery policy general settings.

Server Timeout

The amount of time that passes, in minutes, before the system drops a server from the network map due to inactivity. The default setting is 10080 minutes (one week).

To avoid premature timeout of servers, make sure that the service timeout value is longer than the update interval in the network discovery policy general settings.

Client Application Timeout

The amount of time that passes, in minutes, before the system drops a client from the network map due to inactivity. The default setting is 10080 minutes (one week).

Make sure that the client timeout value is longer than the update interval in the network discovery policy general settings.

Related Topics

Host Limit, on page 2383

Configuring Network Discovery Data Storage

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Edit** (✎) next to **Network Discovery Data Storage Settings**.

Step 4 Update the settings in the Data Storage Settings dialog as described in Network Discovery Data Storage Settings, on page 2659.

Step 5 Click **Save** to save the data storage settings.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Configuring Network Discovery Event Logging

The Event Logging Settings control whether discovery and host input events are logged. If you do not log an event, you cannot retrieve it in event views or use it to trigger correlation rules.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Edit** (✎) next to **Event Logging Settings**.

Step 4 Check or clear the check boxes next to the discovery and host input event types you want to log in the database , described in the *Discovery Events* chapter of the Cisco Secure Firewall Management Center Administration Guide.

Step 5 Click **Save** to save the event logging settings.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Adding Network Discovery OS and Server Identity Sources

In Advanced of the network discovery policy, you can add new active sources or change the priority or timeout settings for existing sources.

Adding a scanner to this page does not add the full integration capabilities that exist for the Nmap scanners, but does allow integration of imported third-party application or scan results.

If you import data from a third-party application or scanner, make sure that you map vulnerabilities from the source to the vulnerabilities detected in your network.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Edit** (✎) next to **OS and Server Identity Sources**.

Step 4 To add a new source, click **Add Source**.

Step 5 Enter a **Name**.

Step 6 Choose the input source **Type** from the drop-down list:

- Choose **Scanner** if you plan to import scan results using the AddScanResult function.
- Choose **Application** if you do not plan to import scan results.

Step 7 To indicate the duration of time that should elapse between the addition of an identity to the network map by this source and the deletion of that identity, choose **Hours**, **Days**, or **Weeks** from the **Timeout** drop-down list and enter the appropriate duration.

Step 8 Optionally:

- To promote a source and cause the operating system and application identities to be used in favor of sources below it in the list, choose the source and click the up arrow.
- To demote a source and cause the operating system and application identities to be used only if there are no identities provided by sources above it in the list, choose the source and click the down arrow.
- To delete a source, click **Delete** (🗑) next to the source.

Step 9 Click **Save** to save the identity source settings.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Related Topics

Mapping Third-Party Vulnerabilities, on page 2600

Troubleshooting Your Network Discovery Strategy

Before you make any changes to the system's default detection capabilities, you should analyze what hosts are not being identified correctly and why, so you can decide what solution to implement.

Are Your Managed Devices Correctly Placed?

If network devices such as load balancers, proxy servers, or NAT devices reside between the managed device and the unidentified or misidentified host, place a managed device closer to the misidentified host rather than using custom fingerprinting. Cisco does not recommend using custom fingerprinting in this scenario.

Do Unidentified Operating Systems Have a Unique TCP Stack?

If the system misidentifies a host, you should investigate why the host is misidentified to help you decide between creating and activating a custom fingerprint or substituting Nmap or host input data for discovery data.

Caution If you encounter misidentified hosts, contact your support representative before creating custom fingerprints.

If a host is running an operating system that is not detected by the system by default and does not share identifying TCP stack characteristics with existing detected operating systems, you should create a custom fingerprint.

For example, if you have a customized version of Linux with a unique TCP stack that the system cannot identify, you would benefit from creating a custom fingerprint, which allows the system to identify the host and continuing monitoring it, rather than using scan results or third-party data, which require you to actively update the data yourself on an ongoing basis.

Note that many open source Linux distributions use the same kernel, and as such, the system identifies them using the Linux kernel name. If you create a custom fingerprint for a Red Hat Linux system, you may see other operating systems (such as Debian Linux, Mandrake Linux, Knoppix, and so on) identified as Red Hat Linux, because the same fingerprint matches multiple Linux distributions.

You should not use a fingerprint in every situation. For example, a modification may have been made to a host's TCP stack so that it resembles or is identical to another operating system. For example, an Apple Mac OS X host is altered, making its fingerprint identical to a Linux 2.4 host, causing the system to identify it as Linux 2.4 instead of Mac OS X. If you create a custom fingerprint for the Mac OS X host, it may cause all legitimate Linux 2.4 hosts to be erroneously identified as Mac OS X hosts. In this case, if Nmap correctly identifies the host, you could schedule regular Nmap scans for that host.

If you import data from a third-party system using host input, you must map the vendor, product, and version strings that the third party uses to describe servers and application protocols to the Cisco definitions for those products. Note that even if you map application data to system vendor and version definitions, imported third-party vulnerabilities are not used for impact assessment for clients or web applications.

The system may reconcile data from multiple sources to determine the current identity for an operating system or application.

For Nmap data, you can schedule regular Nmap scans. For host input data, you can regularly run the Perl script for the import or the command line utility. However, note that active scan data and host input data may not be updated with the frequency of discovery data.

Can the System Identify All Applications?

If a host is correctly identified by the system but has unidentified applications, you can create a user-defined detector to provide the system with port and pattern matching information to help identify the application.

Have You Applied Patches that Fix Vulnerabilities?

If the system correctly identifies a host but does not reflect applied fixes, you can use the host input feature to import patch information. When you import patch information, you must map the fix name to a fix in the database.

Do You Want to Track Third-Party Vulnerabilities?

If you have vulnerability information from a third-party system that you want to use for impact correlation, you can map the third-party vulnerability identifiers for servers and application protocols to vulnerability identifiers in the Cisco database and then import the vulnerabilities using the host input feature. For more information on using the host input feature, see the *Firepower System Host Input API Guide*. Note that even if you map application data to system vendor and version definitions, imported third-party vulnerabilities are not used for impact assessment for clients or web applications.

PART XIV

FlexConfig Policies

- FlexConfig Policies, on page 2667

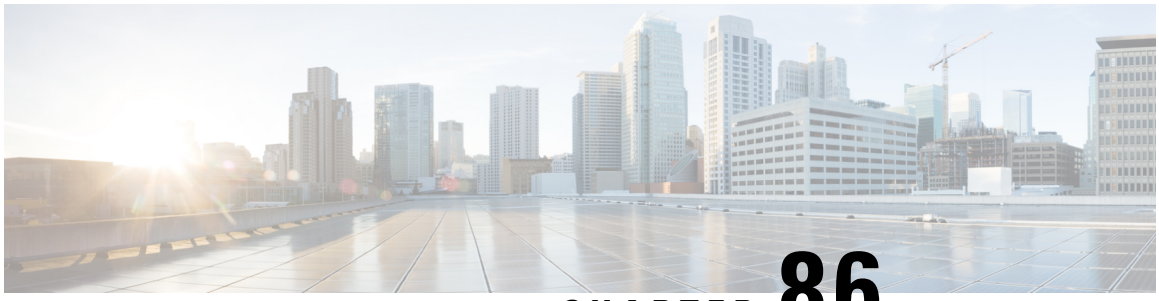

CHAPTER 86

FlexConfig Policies

The following topics describe how to configure and deploy FlexConfig policies.

- FlexConfig Policy Overview, on page 2667
- Requirements and Prerequisites for FlexConfig Policies, on page 2687
- Guidelines and Limitations for FlexConfig, on page 2687
- Customizing Device Configuration with FlexConfig Policies, on page 2688
- Examples for FlexConfig, on page 2701
- Migrating FlexConfig Policies, on page 2708
- History for FlexConfig, on page 2710

FlexConfig Policy Overview

A FlexConfig policy is a container of an ordered list of FlexConfig objects. Each object includes a series of Apache Velocity scripting language commands, ASA software configuration commands, and variables that you define. The contents of each FlexConfig object is essentially a program that generates a sequence of ASA commands that will then be deployed to the assigned devices. This command sequence then configures the related feature on the Firewall Threat Defense device.

Firewall Threat Defense uses ASA configuration commands to implement some features, but not all features. There is no unique set of Firewall Threat Defense configuration commands. Instead, the point of FlexConfig is to allow you to configure features that are not yet directly supported through Firewall Management Center policies and settings.

Caution Cisco **strongly** recommends using FlexConfig policies only if you are an advanced user with a strong ASA background and at your own risk. You may configure any commands that are not prohibited. Enabling features through FlexConfig policies may cause unintended results with other configured features.

You may contact the Cisco Technical Assistance Center for support concerning FlexConfig policies that you have configured. The Cisco Technical Assistance Center does not design or write custom configurations on any customer's behalf. Cisco expresses no guarantees for correct operation or interoperability with other Firepower System features. FlexConfig features may become deprecated at any time. For fully guaranteed feature support, you must wait for the Firewall Management Center support. When in doubt, do not use FlexConfig policies.

Recommended Usage for FlexConfig Policies

There are two main recommended uses for FlexConfig:

- You are converting from ASA to Firewall Threat Defense, and there are compatible features you are using (and need to continue using) that Firewall Management Center does not directly support. In this case, use the **show running-config** command on the ASA to see the configuration for the feature and create your FlexConfig objects to implement it. Experiment with the object's deployment settings (once/every time and append/prepend) to get the right setting. Verify by comparing **show running-config** output on the two devices.
- You are using Firewall Threat Defense but there is a setting or feature that you need to configure, e.g. the Cisco Technical Assistance Center tells you that a particular setting should resolve a specific problem you are encountering. For complicated features, use a lab device to test the FlexConfig and verify that you are getting the expected behavior.

The system includes a set of predefined FlexConfig objects that represent tested configurations. If the feature you need is not represented by these objects, first determine if you can configure an equivalent feature in standard policies. For example, the access control policy includes intrusion detection and prevention, HTTP and other types of protocol inspection, URL filtering, application filtering, and access control, which the ASA implements using separate features. Because many features are not configured using CLI commands, you will not see every policy represented within the output of **show running-config**.

Note At all times, keep in mind that there is not a one-to-one overlap between ASA and Firewall Threat Defense. Do not attempt to completely recreate an ASA configuration on a Firewall Threat Defense device. You must carefully test any feature that you configure using FlexConfig.

CLI Commands in FlexConfig Objects

The Firewall Threat Defense uses ASA configuration commands to configure some features. Although not all ASA features are compatible with the Firewall Threat Defense, there are some features that can work on the Firewall Threat Defense but that you cannot configure in the Firewall Management Center policies. You can use FlexConfig objects to specify the CLI required to configure these features.

If you decide to use FlexConfig to manually configure a feature, you are responsible for knowing and implementing the commands according to the proper syntax. FlexConfig policies do not validate CLI command syntax. For more information about proper syntax and configuring CLI commands, use the ASA documentation as a reference:

- ASA CLI configuration guides explain how to configure a feature. Find the guides at https://www.cisco.com/c/en/us/support/security/adaptive-security-appliance-asa-software/products-installation-and-configuration-guides-list.html
- ASA command references provide additional information sorted by command name. Find the references at https://www.cisco.com/c/en/us/support/security/adaptive-security-appliance-asa-software/products-command-reference-list.html

The following topics explain more about configuration commands.

Determine the ASA Software Version and Current CLI Configuration

Because the system uses ASA software commands to configure some features, you need to determine the current ASA version used in software running on the Firewall Threat Defense device. This version number indicates which ASA CLI configuration guides to use for instructions on configuring a feature. You also should examine the current CLI-based configuration and compare it to the ASA configuration you want to implement.

Keep in mind that any ASA configuration will be very different from a Firewall Threat Defense configuration. Many Firewall Threat Defense policies are configured outside of the CLI, so you cannot see the configuration by looking at the commands. Do not try to create a one-to-one correspondence between an ASA and Firewall Threat Defense configuration.

To view this information, make an SSH connection to the device's management interface and issue the following commands:

- **show version system** and look for the Cisco Adaptive Security Appliance Software Version number. (If you issue the command through the Secure Firewall Management Center CLI tool, omit the **system** keyword.)
- **show running-config** to view the current CLI configuration.
- **show running-config all** to include all the default commands in the current CLI configuration.

You can also issue these commands from within Firewall Management Center using the following procedure.

Procedure

Step 1 Choose **System** (⚙) > **Health** > **Monitor**.

Step 2 Click the name of the device targeted by the FlexConfig policy.

You might need to click the open/close arrow in the **Count** column in the Status table to see any devices.

Step 3 Click **View System and Troubleshoot Details**.

Step 4 Click **Advanced Troubleshooting**.

Step 5 Click **Threat Defense CLI**.

Step 6 Choose the **Device**, then choose **show** as the command, and type **version** or one of the other commands as the parameter.

Step 7 Click **Execute**.

For version, search the output for the Cisco Adaptive Security Appliance Software Version number.

You can select the output and press Ctrl+C, then paste it into a text file for later analysis.

Prohibited CLI Commands

The purpose of FlexConfig is to configure features that are available on ASA devices that you cannot configure on Firewall Threat Defense devices using Firewall Management Center.

Thus, you are prevented from configuring ASA features that have equivalents in Firewall Management Center. The following table lists some of these prohibited command areas.

In addition, some **clear** commands are prohibited because they overlap with managed policies, and can delete part of the configuration for a managed policy.

The FlexConfig object editor prevents you from including prohibited commands in the object.

Prohibited CLI Command	Description
AAA	Configuration blocked.
AAA-Server	Configuration blocked.
Access-list	Advanced ACL, Extended ACL, and Standard ACL are blocked. Ethertype ACL is allowed. You can use standard and extended ACL objects defined in the object manager inside the template as variables.
ARP Inspection	Configuration blocked.
As-path Object	Configuration blocked.
Banner	Configuration blocked.
BGP	Configuration blocked.
Clock	Configuration blocked.
Community-list Object	Configuration blocked.
Copy	Configuration blocked.
Delete	Configuration blocked.
DHCP	Configuration blocked.
Enable Password	Configuration blocked.
Erase	Configuration blocked.
Fragment Setting	Blocked, except for **fragment reassembly**.
Fsck	Configuration blocked.
HTTP	Configuration blocked.
ICMP	Configuration blocked.
Interface	Only **nameif**, **mode**, **shutdown**, **ip address** and **mac-address** commands are blocked.
Multicast Routing	Configuration blocked.
NAT	Configuration blocked.
Network Object/Object-group	Network object creation in the FlexConfig object is blocked, but you can use network objects and groups defined in the object manager inside the template as variables.

Prohibited CLI Command	Description
NTP	Configuration blocked.
OSPF/OSPFv3	Configuration blocked.
pager	Configuration blocked.
Password Encryption	Configuration blocked.
Policy-list Object	Configuration blocked.
Prefix-list Object	Configuration blocked.
Reload	You cannot schedule reloads. The system does not use the **reload** command to restart the system, it uses the **reboot** command.
RIP	Configuration blocked.
Route-Map Object	Route-map object creation in the FlexConfig object is blocked, but you can use route map objects defined in the object manager inside the template as variables.
Service Object/Object-group	Service object creation in the FlexConfig object is blocked, but you can use port objects defined in the object manager inside the template as variables.
SNMP	Configuration blocked.
SSH	Configuration blocked.
Static Route	Configuration blocked.
Syslog	Configuration blocked.
Time Synchronization	Configuration blocked.
Timeout	Configuration blocked.
VPN	Configuration blocked.

Template Scripts

You can use scripting language to control processing within a FlexConfig object. Scripting language instructions are a subset of commands supported in the Apache Velocity 1.3.1 template engine, a Java-based scripting language that supports looping, if/else statements, and variables.

To learn how to use the scripting language, see the *Velocity Developer Guide* at http://velocity.apache.org/engine/devel/developer-guide.html.

FlexConfig Variables

You can use variables in a FlexConfig object in cases where part of a command or processing instruction depends on runtime information rather than static information. During deployment, the variables are replaced with strings obtained from other configurations for the device based on the type of variable:

- Policy object variables are replaced with strings obtained from objects defined in Firewall Management Center.
- System variables are replaced with information obtained from the device itself or from policies configured for it.
- Processing variables are loaded with the contents of policy object or system variables as scripting commands are processed. For example, in a loop, you iteratively load one value from a policy object or system variable into a processing variable, then use the processing variable to form a command string or perform some other action. These processing variables do not show up in the Variables list within a FlexConfig object. Also, you do not add them using the **Insert** menu in the FlexConfig object editor.
- Secret key variables are replaced with the single string defined for the variable within the FlexConfig object.

Variables start with the $ character, except for secret keys, which start with the @ character. For example, $ifname is a policy object variable in the following command, whereas @keyname is a secret key.

```
interface $ifname
key @keyname
```

Note The first time you insert a policy object or system variable, you must do so through the **Insert** menu in the FlexConfig object editor. This action adds the variable to the **Variables** list at the bottom of the FlexConfig object editor. But you must type in the variable string on subsequent uses, even when using system variables. If you are adding a processing variable, which does not have an object or system variable assignment, do not use the **Insert** menu. If you are adding a secret key, always use the **Insert** menu. Secret key variables do not show up in the Variables list.

Whether a variable is resolved as a single string, a list of strings, or a table of values depends on the type of policy object or system variable you assign to the variable. (Secret keys always resolve to a single string.) You must understand what will be returned in order to process the variables correctly.

The following topics explain the various types of variable and how to process them.

How to Process Variables

At runtime, a variable can resolve to a single string, a list of strings of the same type, a list of strings of different types, or a table of named values. In addition, variables that resolve to multiple values can be of determinate or indeterminate length. You must understand what will be returned in order to process the values correctly.

Following are the main possibilities.

Single Value Variables

If a variable always resolves to a single string, use the variable directly without modification in the FlexConfig script.

For example, the predefined text variable tcpMssBytes always resolves to a single value (which must be numeric). The **Sysopt_basic** FlexConfig then uses an if/then/else structure to set the maximum segment size based on the value of another single-value text variable, tcpMssMinimum:

```
#if($tcpMssMinimum == "true")
  sysopt connection tcpmss minimum $tcpMssBytes
#else
  sysopt connection tcpmss $tcpMssBytes
#end
```

In this example, you would use the **Insert** menu in the FlexConfig object editor to add the first use of $tcpMssBytes, but you would type in the variable directly on the #else line.

Secret key variables are a special type of single value variable. For secret keys, you always use the **Insert** menu to add the variable, even for second and subsequent uses. These variables do not show up in the Variables list within the FlexConfig object.

Note Policy object variables for network objects also equate to a single IP address specification, either a host address, network address, or address range. However, in this case, you must know what type of address to expect, because the ASA commands require specific address types. For example, if a command requires a host address, using a network object variable that points to an object that contains a network address will result in an error during deployment.

Multiple Value Variables, All Values Are the Same Type

Several policy object and system variables resolve to multiple values of the same type. For example, an object variable that points to a network object group resolves to a list of the IP addresses within the group. Similarly, the system variable $SYS_FW_INTERFACE_NAME_LIST resolves to a list of interface names.

You can also create text objects for multiple values of the same type. For example, the predefined text object enableInspectProtocolList can contain more than one protocol name.

Multiple value variables that resolve to a list of items of the same type are frequently of indeterminate length. For example, you cannot know beforehand how many interfaces on a device are named, as users can configure or unconfigure interfaces at any time.

Thus, you would typically use a loop to process multiple value variables of the same type. For example, the predefined FlexConfig **Default_Inspection_Protocol_Enable** uses a #foreach loop to go through the enableInspectProtocolList object and process each value.

```
policy-map global_policy
    class inspection_default
    #foreach ( $protocol in $enableInspectProtocolList)
    inspect $protocol
    #end
```

In this example, the script assigns each value in turn to the $protocol variable, which is then used in an ASA **inspect** command to enable the inspection engine for that protocol. In this case, you simply type in $protocol as a variable name. You do not use the **Insert** menu to add it, because you are not assigning an object or system value to the variable. However, you must use the **Insert** menu to add $enableInspectProtocolList.

The system loops through the code between #foreach and #end until there are no values remaining in $enableInspectProtocolList.

Multiple Value Variables, Values Are Different Types

You can create multiple value text objects, but have each value serve a different purpose. For example, the predefined **netflow_Destination** text object should have 3 values, in order, interface name, destination IP address, and UDP port number.

Objects defined in this way should have a determinate number of values. Otherwise, they would be hard to process.

Use the get method to process these objects. Type **.get(*n*)** at the end of the object name, replacing *n* with an index into the object. Start counting at 0, even though the text object lists its values starting at 1.

For example, the Netflow_Add_Destination object uses the following line to add the 3 values from netflow_Destination to the ASA **flow-export** command.

```
flow-export destination $netflow_Destination.get(0) $netflow_Destination.get(1)
$netflow_Destination.get(2)
```

In this example, you would use the **Insert** menu in the FlexConfig object editor to add the first use of $netflow_Destination, and then add **.get(0)**. But you would type in the variable directly for the **$netflow_Destination.get(1)** and **$netflow_Destination.get(2)** specifications.

Multiple Value Variables that Resolve to a Table of Values

Some system variables return a table of values. These variables include MAP in their name, for example, $SYS_FTD_ROUTED_INTF_MAP_LIST. The routed interface map returns data that looks like the following (line returns added for clarity):

```
[{intf_hardwarare_id=GigabitEthernet0/0, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=255.255.255.0,
intf_ip_addr_v4=10.100.10.1, intf_ipv6_link_local_address=,
intf_logical_name=outside},

{intf_hardwarare_id=GigabitEthernet0/1, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=255.255.255.0,
intf_ip_addr_v4=10.100.11.1, intf_ipv6_link_local_address=,
intf_logical_name=inside},

{intf_hardwarare_id=GigabitEthernet0/2, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=, intf_ip_addr_v4=,
intf_ipv6_link_local_address=, intf_logical_name=},

{intf_hardwarare_id=Management0/0, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=, intf_ip_addr_v4=,
intf_ipv6_link_local_address=, intf_logical_name=management}]
```

In the above example, information is returned for 4 interfaces. Each interface includes a table of named values. For example, intf_hardwarare_id is the name of the interface hardware name property, and returns strings such as GigabitEthernet0/0.

This type of variable is typically of indeterminate length, so you need to use looping to process the values. But you also need to add the property name to the variable name to indicate which value to retrieve.

For example, IS-IS configuration requires that you add the ASA **isis** command to an interface that has a logical name in interface configuration mode. However, you enter that mode using the interface's hardware name. Thus, you need to identify which interfaces have logical names, then configure just those interfaces using their hardware names. The ISIS_Interface_Configuration predefined FlexConfig does this using an if/then

structure nested in a loop. In the following code, you can see that the #foreach scripting command loads each interface map into the $intf variable, then the #if statement keys off the intf_logical_name value in the map ($intf.intf_logical_name), and if the value is in the list defined in the isisIntfList predefined text variable, enters the interface command using the intf_hardwarare_id value ($intf.intf_hardwarare_id). You would need to edit the isisIntfList variable to add the names of the interfaces on which to configure IS-IS.

```
#foreach ($intf in $SYS_FTD_ROUTED_INTF_MAP_LIST)
 #if ($isIsIntfList.contains($intf.intf_logical_name))
  interface $intf.intf_hardwarare_id
   isis
   #if ($isIsAddressFamily.contains("ipv6"))
   ipv6 router isis
   #end
 #end
#end
```

How to See What a Variable Will Return for a Device

An easy way to evaluate what a variable will return is to create a simple FlexConfig object that does nothing more than process an annotated list of variables. Then, you can assign it to a FlexConfig policy, assign the policy to a device, save the policy, then preview the configuration for that device. The preview will show the resolved values. You can select the preview text, press Ctrl+C, then paste the output into a text file for analysis.

Note Do not deploy this FlexConfig to the device, however, because it will not contain any valid configuration commands. You would get deployment errors. After obtaining the preview, delete the FlexConfig object from the FlexConfig policy and save the policy.

For example, you could construct the following FlexConfig object:

```
Following is a network object group variable for the
IPv4-Private-All-RFC1918 object:

$IPv4_Private_addresses

Following is the system variable SYS_FW_MANAGEMENT_IP:

$SYS_FW_MANAGEMENT_IP

Following is the system variable SYS_FW_ENABLED_INSPECT_PROTOCOL_LIST:

$SYS_FW_ENABLED_INSPECT_PROTOCOL_LIST

Following is the system variable SYS_FTD_ROUTED_INTF_MAP_LIST:

$SYS_FTD_ROUTED_INTF_MAP_LIST

Following is the system variable SYS_FW_INTERFACE_NAME_LIST:

$SYS_FW_INTERFACE_NAME_LIST
```

The preview of this object might look like the following (line returns added for clarity):

```
###Flex-config Prepended CLI ###
```

```
###CLI generated from managed features ###

###Flex-config Appended CLI ###
Following is an network object group variable for the
IPv4-Private-All-RFC1918 object:

[10.0.0.0, 172.16.0.0, 192.168.0.0]

Following is the system variable SYS_FW_MANAGEMENT_IP:

192.168.0.171

Following is the system variable SYS_FW_ENABLED_INSPECT_PROTOCOL_LIST:

[dns, ftp, h323 h225, h323 ras, rsh, rtsp, sqlnet, skinny, sunrpc,
xdmcp, sip, netbios, tftp, icmp, icmp error, ip-options]

Following is the system variable SYS_FTD_ROUTED_INTF_MAP_LIST:

[{intf_hardwarare_id=GigabitEthernet0/0, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=255.255.255.0,
intf_ip_addr_v4=10.100.10.1, intf_ipv6_link_local_address=,
intf_logical_name=outside},

{intf_hardwarare_id=GigabitEthernet0/1, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=255.255.255.0,
intf_ip_addr_v4=10.100.11.1, intf_ipv6_link_local_address=,
intf_logical_name=inside},

{intf_hardwarare_id=GigabitEthernet0/2, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=, intf_ip_addr_v4=,
intf_ipv6_link_local_address=, intf_logical_name=},

{intf_hardwarare_id=Management0/0, intf_ipv6_eui64_addresses=[],
intf_ipv6_prefix_addresses=[], intf_subnet_mask_v4=, intf_ip_addr_v4=,
intf_ipv6_link_local_address=, intf_logical_name=management}]

Following is the system variable SYS_FW_INTERFACE_NAME_LIST:

[outside, inside, management]
```

FlexConfig Policy Object Variables

A policy object variable is associated with a specific policy object configured in the Object Manager. When you insert a policy object variable in a FlexConfig object, you give the variable a name and select the object associated with it.

Although you can give the variable the exact same name as the associated object, the variable itself is not the same thing as the associated object. You must use the **Insert** > **Insert Policy Object** > ***Object Type*** menu in the FlexConfig object editor to add the variable for the first time to the script in the FlexConfig to establish the association with the object. Simply typing in the name of the object preceded by a $ sign does not create a policy object variable.

You can create variables to point to the following types of object. Ensure that you create the right type of object for each variable. To create objects, go to the **Objects** > **Object Management** page.

- **Text Objects**—For text strings, which can include IP addresses, numbers, and other free-form text such as interface or zone names. Select **FlexConfig** > **Text Object** from the table of contents, then click **Add Text Object**. You can configure these objects to contain a single value or multiple values. These objects

are highly flexible and built specifically for use within FlexConfig objects. For detailed information, see Configure FlexConfig Text Objects, on page 2694.

- **Network**—For IP addresses. You can use network objects or groups. Select **Network** from the table of contents, then select **Add Network** > **Add Object** or **Add Group**. If you use a group object, the variable returns a list of each IP address specification within the group. Addresses can be host, network, or address ranges, depending on the object contents. See Network, on page 1401.
- **Security Zones**—For interfaces within a security zone or interface group. Select **Interface** from the table of contents, then select **Add** > **Security Zone** or **Interface Group**. A security zone variable returns a list of the interfaces within that zone or group for the device being configured. See Interface, on page 1398.
- **Standard ACL Object**—For standard access control lists. A standard ACL variable returns the name of the standard ACL object. Select **Access List** > **Standard** from the table of contents, then click **Add Standard Access List Object**. See Access List, on page 1369.
- **Extended ACL Object**—For extended access control lists. An extended ACL variable returns the name of the extended ACL object. Select **Access List** > **Extended** from the table of contents, then click **Add Extended Access List Object**. See Access List, on page 1369.
- **Route Map**—For route map objects. A route map variable returns the name of the route map object. Select **Route Map** from the table of contents, then click **Add Route Map**. See Route Map, on page 1426.

FlexConfig System Variables

System variables are replaced with information obtained from the device itself or from policies configured for it.

You must use the **Insert** > **Insert System Variable** > ***Variable Name*** menu in the FlexConfig object editor to add the variable for the first time to the script in the FlexConfig to establish the association with the system variable. Simply typing in the name of the system variable preceded by a $ sign does not create a system variable within the context of the FlexConfig object.

The following table explains the available system variables. Before using a variable, examine what is typically returned for the variable; see How to See What a Variable Will Return for a Device, on page 2675.

Name	Description
SYS_FW_OS_MODE	The operating system mode of the device. Possible values are ROUTED or TRANSPARENT.
SYS_FW_OS_MULTIPLICITY	Whether the device is running in single or multiple context mode. Possible values are SINGLE, MULTI, or NOT_APPLICABLE.
SYS_FW_MANAGEMENT_IP	The management IP address of the device
SYS_FW_HOST_NAME	The device hostname
SYS_FTD_INTF_POLICY_MAP	A map with interface name as key and policy-map as value. This variable returns nothing if there are no interface-based service policies defined on the device.
SYS_FW_ENABLED_INSPECT_PROTOCOL_LIST	The list of protocols for which inspection is enabled.

Name	Description
SYS_FTD_ROUTED_INTF_MAP_LIST	A list of routed interface maps on the device. Each map includes a set of named values related to routed interface configuration.
SYS_FTD_SWITCHED_INTF_MAP_LIST	A list of switched interface maps on the device. Each map includes a set of named values related to switched interface configuration.
SYS_FTD_INLINE_INTF_MAP_LIST	A list of inline interface maps on the device. Each map includes a set of named values related to inline set interface configuration.
SYS_FTD_PASSIVE_INTF_MAP_LIST	A list of passive interface maps on the device. Each map includes a set of named values related to passive interface configuration.
SYS_FTD_INTF_BVI_MAP_LIST	A list of Bridge Virtual Interface maps on the device. Each map includes a set of named values related to BVI configuration.
SYS_FW_INTERFACE_HARDWARE_ID_LIST	A list of the hardware names for interfaces on the device, such as GigabitEthernet0/0.
SYS_FW_INTERFACE_NAME_LIST	A list of logical names for interfaces on the device, such as inside.
SYS_FW_INLINE_INTERFACE_NAME_LIST	A list of logical names for interfaces configured as passive or ERSPAN passive.
SYS_FW_NON_INLINE_INTERFACE_NAME_LIST	A list of logical names for interfaces that are not part of inline sets, such as all routed interfaces.

Predefined FlexConfig Objects

The predefined FlexConfig objects provide tested configurations for select features. Use these objects if you need to configure these features, which otherwise cannot be configured using the Firewall Management Center.

The following table lists the available objects. Make note of the associated text objects. You must edit these text objects to customize the behavior of the predefined FlexConfig object. The text objects make it possible for you to customize the configuration using the IP addresses and other attributes required by your network and device.

If you need to modify a predefined FlexConfig object, copy the object, make changes to the copy, and save it with a new name. You cannot directly edit a predefined FlexConfig object.

Although you might be able to configure other ASA-based features using FlexConfig, the configuration of those features has not been tested. If an ASA feature overlaps with something that you can configure in Firewall Management Center policies, do not attempt to configure it through FlexConfig.

For example, Snort inspection includes the HTTP protocol, so do not enable ASA-style HTTP inspection. (In fact, you cannot add **http** to the enableInspectProtocolList object. In this case, you are prevented from misconfiguring your device.) Instead, configure the access control policy to perform application or URL filtering, as needed, to implement your HTTP inspection requirements.

Table 144: Predefined FlexConfig Objects

FlexConfig Object Name	Description	Associated Text Objects
Default_Inspection_Protocol_Disable	Disables protocols in the global_policy default policy map.	disableInspectProtocolList
Default_Inspection_Protocol_Enable	Enables protocols in the global_policy default policy map.	enableInspectProtocolList
Inspect_IPv6_Configure	Configures IPv6 inspection in the global_policy policy map, logging and dropping traffic based on IPv6 header contents.	IPv6RoutingHeaderDropLogList, IPv6RoutingHeaderLogList, IPv6RoutingHeaderDropList.
Inspect_IPv6_UnConfigure	Clears and disables IPv6 inspection.	—
ISIS_Configure	Configures global parameters for IS-IS routing.	isIsNet, isIsAddressFamily, isISType
ISIS_Interface_Configuration	Interface level IS-IS configuration.	isIsAddressFamily, IsIsIntfList Also uses the system variable SYS_FTD_ROUTED_INTF_MAP_LIST
ISIS_Unconfigure	Clears the IS-IS router configuration on the device.	—
ISIS_Unconfigure_All	Clears the IS-IS router configuration from the device, including the router assignment from the device interface.	—
NGFW_TCP_NORMALIZATION	Modifies the default TCP normalization configuration.	—
Policy_Based_Routing	To use this example configuration, copy it, modify the interface name, and use the r-map-object text object to identify a route map object in the object manager.	—
Policy_Based_Routing_Clear	Clears Policy Based Routing configurations from the device.	—
Sysopt_AAA_radius	Ignores the authentication key in RADIUS accounting responses.	—
Sysopt_AAA_radius_negate	Negates the Sysopt_AAA_radius configuration.	—
Sysopt_basic	Configures sysopt wait time , maximum segment size for TCP packets, and detailed traffic statistics.	tcpMssMinimum, tcpMssBytes

FlexConfig Object Name	Description	Associated Text Objects
Sysopt_basic_negate	Clears sysopt_basic detailed traffic statistics, wait time, and TCP maximum segment size.	—
Sysopt_clear_all	Clears all sysopt configurations from the device.	—
Sysopt_noproxyarp	Configures noproxy-arp CLIs.	Uses system variable SYS_FW_NON_INLINE_INTF_NAME_LIST
Sysopt_noproxyarp_negate	Clears Sysopt_noproxyarp configurations.	Uses system variable SYS_FW_NON_INLINE_INTF_NAME_LIST
Sysopt_Preserve_Vpn_Flow	Configures syopt preserve VPN flow.	—
Sysopt_Preserve_Vpn_Flow_negate	Clears the Sysopt_Preserve_Vpn_Flow configuration.	—
Sysopt_Reclassify_Vpn	Configures sysopt reclassify vpn.	—
Sysopt_Reclassify_Vpn_Negate	Negates sysopt reclassify vpn.	—
Threat_Detection_Clear	Clear the threat detection TCP Intercept configuration.	—
Threat_Detection_Configure	Configure threat detection statistics for attacks intercepted by TCP Intercept.	threat_detection_statistics
Wccp_Configure	This template provides an example for configuring WCCP.	isServiceIdentifier, serviceIdentifier, wccpPassword
Wccp_Configure_Clear	Clears WCCP configurations.	—

Deprecated FlexConfig Objects

The following table lists objects that configure features you can now configure natively in the GUI. Discontinue using these objects at the earliest convenience.

Table 145: Deprecated Predefined FlexConfig Objects

Deprecated Version	FlexConfig Object	Description	Now Configure In
7.3	DHCPv6_Prefix_Delegation_Configure	Configure one outside (Prefix Delegation client) and one inside interface (recipient of delegated prefix) for IPv6 prefix delegation. To use this template, copy it and modify the variables. Associated text objects: pdoutside, pdinside Also uses the system variable SYS_FTD_ROUTED_INTF_MAP_LIST	Interface IPv6 Settings.
7.3	DHCPv6_Prefix_Delegation_UnConfigure	Removes the DHCPv6 prefix delegation configuration.	Interface IPv6 Settings.
6.3	Default_DNS_Configure	Configure the Default DNS group, which defines the DNS servers that can be used when resolving fully-qualified domain names on the data interfaces. Associated text objects: defaultDNSNameServerList, defaultDNSParameters	Platform settings.
6.3	DNS_Configure	Configure DNS servers in a non-default DNS server group. Copy the object to change the name of the group.	**DNS Server Group** in the object manager.
6.3	DNS_UnConfigure	Removes the DNS server configuration performed by Default_DNS_Configure and DNS_Configure. Copy the object to change the DNS server group names if you altered DNS_Configure.	**DNS Server Group** in the object manager.
7.2	Eigrp_Configure	Configures EIGRP routing next-hop, auto-summary, router-id, eigrp-stub. Associated text objects: eigrpAS, eigrpNetworks, eigrpDisableAutoSummary, eigrpRouterId, eigrpStubReceiveOnly, eigrpStubRedistributed, eigrpStubConnected, eigrpStubStatic, eigrpStubSummary	For all EIGRP objects, see EIGRP, on page 1267. The system does allow you to deploy post-upgrade, but also warns you to redo your EIGRP configurations. To help you with this process, we provide a command-line migration tool.

Deprecated Version	FlexConfig Object	Description	Now Configure In
7.2	Eigrp_Interface_Configure	Configures EIGRP interface authentication mode, authentication key, hello interval, hold time, split horizon. Associated text objects: eigrpIntfList, eigrpAS, eigrpAuthKey, eigrpAuthKeyId, eigrpHelloInterval, eigrpHoldTime, eigrpDisableSplitHorizon Also uses the system variable SYS_FTD_ROUTED_INTF_MAP_LIST	
7.2	Eigrp_Unconfigure	Clears EIGRP configuration for an autonomous system from the device.	
7.2	Eigrp_Unconfigure_all	Clears all EIGRP configurations.	
7.4	Netflow_Add_Destination	Creates and configures a Netflow export destination. Associated text objects: Netflow_Destinations, netflow_Event_Types	Platform settings.
7.4	Netflow_Clear_Parameters	Restores Netflow export global default settings.	Platform settings.
7.4	Netflow_Delete_Destination	Deletes a Netflow export destination. Associated text objects: Netflow_Destinations, netflow_Event_Types	Platform settings.
7.4	Netflow_Set_Parameters	Sets global parameters for Netflow export. Associated text objects: netflow_Parameters	Platform settings.
6.3	TCP_Embryonic_Conn_Limit	Configures embryonic connection limits to protect against SYN Flood Denial of Service (DoS) attacks. Associated text objects: tcp_conn_misc, tcp_conn_limit	Service policy.

Deprecated Version	FlexConfig Object	Description	Now Configure In
6.3	TCP_Embryonic_Conn_Timeout	Configures embryonic connection timeouts to protect against SYN Flood Denial of Service (DoS) attacks. Associated text objects: tcp_conn_misc, tcp_conn_timeout	Service policy.
7.2	VxLAN_Clear_Nve	Removes the NVE 1 configured when VxLAN_Configure_Port_And_Nve is used from the device.	For all VxLAN objects, see Configure VXLAN Interfaces, on page 838. If you configured VXLAN interfaces with FlexConfig in a previous version, they continue to work. In fact, FlexConfig takes precedence in this case—if you redo your VXLAN configurations in the web interface, remove the FlexConfig settings.
7.2	VxLAN_Clear_Nve_Only	Clears the NVE configured on the interface when deployed.	
7.2	VxLAN_Configure_Port_And_Nve	Configures VLAN port and NVE 1. Associated text objects: vxlan_Port_And_Nve	
7.2	VxLAN_Make_Nve_Only	Sets an interface for NVE only. Associated text objects: vxlan_Nve_Only Also uses system variables SYS_FTD_ROUTED_MAP_LIST and SYS_FTD_SWITCHED_INTF_MAP_LIST	
7.2	VxLAN_Make_Vni	Creates a VNI interface. After deploying this you have to unregister and re-register the device to properly discover the VNI interface. Associated text objects: vxlan_Vni	

Predefined Text Objects

There are several predefined text objects. These objects are associated with variables used in the predefined FlexConfig objects. In most cases, you must edit these objects to add values if you use the associated FlexConfig object, or you will see errors during deployment. Although some of these objects contain default values, others are empty.

For information on editing text objects, see Configure FlexConfig Text Objects, on page 2694.

Name	Description	Associated FlexConfig Object
defaultDNSNameServerList (Deprecated.)	The DNS server IP address to configure in the Default DNS group. Starting with version 6.3, configure DNS for the data interfaces in the Threat Defense Platform Settings policy.	Default_DNS_Configure
defaultDNSParameters (Deprecated.)	The parameters to control DNS behavior for the default DNS server group. The object contains separate entries, in order, for retries, timeout, expire-entry-timer, poll-timer, domain-name. Starting with version 6.3, configure DNS for the data interfaces in the Threat Defense Platform Settings policy.	Default_DNS_Configure
disableInspectProtocolList	Disables protocols in the default policy map (global_policy).	Disable_Default_Inspection_Protocol
dnsNameServerList	The DNS server IP address to configure in a user-defined DNS group.	DNS_Configure
dnsParameters	The parameters to control DNS behavior for a non-default DNS server group. The object contains separate entries, in order, for retries, timeout, domain-name, name-server-interface.	DNS_Configure
enableInspectProtocolList	Enables protocols in the default policy map (global_policy). You are prevented from adding protocols whose inspection conflicts with Snort inspection.	Enable_Default_Inspection_Protocol
IPv6RoutingHeaderDropList	The list of IPv6 routing header types that you want to disallow. IPv6 inspection drops packets that contain these headers without logging the drop.	Inspect_IPv6_Configure
IPv6RoutingHeaderDropLogList	The list of IPv6 routing header types that you want to disallow and log. IPv6 inspection drops packets that contain these headers and sends a syslog message about the drop.	Inspect_IPv6_Configure
IPv6RoutingHeaderLogList	The list of IPv6 routing header types that you want to allow but log. IPv6 inspection allows packets that contain these headers, but sends a syslog message about the existence of the header.	Inspect_IPv6_Configure

Name	Description	Associated FlexConfig Object
isIsAddressFamily	The IPv4 or IPv6 address family.	ISIS_Configure ISIS_Interface_Configuration
IsIsIntfList	List of logical interface names.	ISIS_Interface_Configuration
isIsISType	IS Type (level-1, level-2-only or level-1-2).	ISIS_Configure
isIsNet	Network entity.	ISIS_Configure
isServiceIdentifier	When false, uses the standard **web-cache** service identifier.	Wccp_Configure
netflow_Destination	Defines a single NetFlow export destination's interface, destination, and UDP port number.	Netflow_Add_Destination
netflow_Event_Types	Defines the types of events to be exported for a destination as any subset of: **all**, **flow-create**, **flow-defined**, **flow-teardown**, **flow-update**.	Netflow_Add_Destination
netflow_Parameters	Provides the NetFlow export global settings: active refresh interval (number of minutes between flow update events), delay (flow create delay in seconds; default 0 = command will not appear), and template time-out rate in minutes.	Netflow_Set_Parameters
PrefixDelegationInside	Configures the inside interface for DHCPv6 prefix delegation. The object includes multiple entries, in order, interface name, IPv6 suffix with prefix length, and prefix pool name.	None, but could be used with a copy of DHCPv6_Prefix_Delegation_Configure.
PrefixDelegationOutside	Configure the outside DHCPv6 prefix delegation client. The object includes multiple entries, in order, interface name and IPv6 prefix length	None, but could be used with a copy of DHCPv6_Prefix_Delegation_Configure.
serviceIdentifier	Dynamic WCCP service identifier number.	Wccp_Configure
tcp_conn_limit (Deprecated.)	Parameters used for configuring the TCP embryonic connection limits. Starting with version 6.3, configure these features in the Threat Defense Service Policy, which you can find on the Advanced tab of the access control policy assigned to the device.	TCP_Embryonic_Conn_Limit

Name	Description	Associated FlexConfig Object
tcp_conn_misc (Deprecated.)	Parameters used for configuring the TCP embryonic connection settings. Starting with version 6.3, configure these features in the Threat Defense Service Policy, which you can find on the Advanced tab of the access control policy assigned to the device.	TCP_Embryonic_Conn_Limit, TCP_Embryonic_Conn_Timeout
tcp_conn_timeout (Deprecated.)	Parameters used for configuring the TCP embryonic connection timeouts. Starting with version 6.3, configure these features in the Threat Defense Service Policy, which you can find on the Advanced tab of the access control policy assigned to the device.	TCP_Embryonic_Conn_Timeout
tcpMssBytes	Maximum segment size in bytes.	Sysopt_basic
tcpMssMinimum	Checks whether to set maximum segment size (MSS), which is set only if this flag is true.	Sysopt_basic
threat_detection_statistics	Parameters used for threat detection statistics for TCP Intercept.	Threat_Detection_Configure
vxlan_Nve_Only	Parameters for configuring NVE-only on interface: • logical name of interface • IPv4 address (optional for routed interface) • IPv4 netmask (optional for routed interface)	VxLAN_Make_Nve_Only
vxlan_Port_And_Nve	Parameters used for configuring ports and NVE for VXLAN: • vxlan port • source interface (logical name) • type (peer or mcast) • Peer IP Address or default-mcast-group	VxLAN_Configure_Port_And_Nve

Name	Description	Associated FlexConfig Object
vxlan_Vni	Parameters used for creating VNI: • Interface number (1-10000) • segment-id (1-16777215) • nameif (Logical Name of the interface) • type (routed or transparent) • IP address (used in case of routed mode device) or bridge-group number (used in case of transparent mode device) • netmask (If device is in routed mode) or unused	VxLAN_Make_Vni
wccpPassword	WCCP password.	Wccp_Configure

Requirements and Prerequisites for FlexConfig Policies

Model Support

Threat Defense

Supported Domains

Any

User Roles

Admin

Guidelines and Limitations for FlexConfig

- If you make a mistake in the FlexConfig policy, the system will roll back all changes included in the deployment attempt that includes the failed FlexConfig. Because rollback due to a failed deployment includes clearing the configuration, this can be disruptive to your network. Consider timing deployments that include FlexConfig changes to non-business hours. Also, consider isolation the deployment so it includes just FlexConfig changes, and no other policy updates.
- When you use the VxLAN_Make_VNI object, you must deploy the same FlexConfig to all units in a cluster or high availability pair before you form the cluster or high availability pair. The Management Center requires the VXLAN interfaces to match on all devices before forming the cluster or high availability pair.

- If you configure any service that applies to connections, such as SIP inspection, go to the device CLI and enter the **clear conn** command to clear connections. When the connections are rebuilt, the new configuration is applied to the sessions.

Customizing Device Configuration with FlexConfig Policies

Use FlexConfig policies to customize the configuration of a Firewall Threat Defense device.

Before using FlexConfig, try to configure all the policies and settings you need using the other features in Firewall Management Center. FlexConfig is a method of last resort to configure ASA-based features that are compatible with Firewall Threat Defense but which are not otherwise configurable in Firewall Management Center.

Following is the end-to-end procedure for configuring and deploying a FlexConfig policy.

Procedure

Step 1 Determine the CLI command sequence that you want to configure.

If you have a functioning configuration on an ASA device, use **show running-config** to get the sequence of commands that you need. Make adjustments to items such as interface names and IP addresses as needed.

If this is for a new feature, it is best to try to implement it on an ASA device in a lab setting to verify that you have the correct command sequence.

For more information, see the following topics:

- Recommended Usage for FlexConfig Policies, on page 2668
- CLI Commands in FlexConfig Objects, on page 2668

Step 2 Choose **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object**.

Examine the predefined FlexConfig objects to determine if any will be able to generate the commands you need. Click **View** ( ) to see the object contents. If an existing object is close to what you want, start by making a copy of the object, and then edit the copy. See Predefined FlexConfig Objects, on page 2678.

Examining the objects will also give you an idea of the structure, command syntax, and expected sequencing for a FlexConfig object.

Note
If you find any objects that you will use, either directly or as copies, examine the Variables list at the bottom of the object. Make note of the variable names, except those in all capitals that start with SYS, which are system variables. These variables are text objects that you will probably need to edit and define overrides for, especially if the default value column shows the object has no value.

Step 3 If you need to create your own FlexConfig objects, determine what variables you will need and create the associated objects.

The CLI you need to deploy might contain IP addresses, interface names, port numbers, and other parameters that you might want to adjust over time. These are best isolated into variables, which point to objects that

contain the necessary values. You might also need variables for strings that are part of the configuration but which might change over time.

Also, determine if you need different values for each device to which you will assign the policy. For example, you might want to configure the feature on three devices, but you might need to specify a different interface name or IP address on a given command for each of these devices. If you need to customize the object for each device, ensure that you enable overrides when creating the object, and then define the override values per device.

See the following topics for an explanation of the various types of variables and how to configure the related objects when necessary.

- FlexConfig Variables, on page 2672
- FlexConfig Policy Object Variables, on page 2676
- FlexConfig System Variables, on page 2677
- Configure FlexConfig Text Objects, on page 2694

Step 4 If you are using the predefined FlexConfig objects, edit the text objects used as variables.

See Configure FlexConfig Text Objects, on page 2694.

Step 5 (If necessary.) Configure FlexConfig Objects, on page 2690.

You need to create objects only if the predefined objects cannot do the job.

Step 6 Configure the FlexConfig Policy, on page 2695.

Step 7 Set Target Devices for a FlexConfig Policy, on page 2696.

You can also assign the policy to devices when you create the policy. The policy must have at least one assigned device before you can preview it.

Step 8 Preview the FlexConfig Policy, on page 2697.

You must save changes before you can preview the policy.

Verify that the generated commands are the ones intended, and that all variables are resolving correctly.

Step 9 Choose **Deploy** > **Deployment** in the menu bar.

Step 10 Select the devices assigned to the policy, and click **Deploy**.

Wait for deployment to complete.

Step 11 Verify the Deployed Configuration, on page 2698.

Step 12 (If necessary.) Remove Features Configured Using FlexConfig, on page 2699.

Unlike other types of policy, simply unassigning a FlexConfig from a device might not remove the related configuration. If you want to remove a FlexConfig-generated configuration, you follow the cited procedure.

If you are removing a Feature because it is now directly supported by the product, see also Convert from FlexConfig to Managed Feature, on page 2701.

Configure FlexConfig Objects

Use FlexConfig objects to define a configuration to be deployed to a device. Each FlexConfig policy is composed of a list of FlexConfig objects, so the objects are essentially code modules composed of Apache Velocity scripting commands, ASA software configuration commands, and variables.

There are several predefined FlexConfig objects that you can use directly, or you can make copies if you need to edit them. You can also create your own objects from scratch. A FlexConfig object's content can range from a single simple command string to elaborate CLI command structures that use variables and scripting commands to deploy commands whose content can differ from device to device or deployment to deployment.

You can also create FlexConfig policy objects when defining FlexConfig policies.

Before you begin

Keep the following in mind:

- FlexConfig objects translate into commands that are then deployed to the device. These commands are already issued in global configuration mode. Therefore, do not include the **enable** and **configure terminal** commands as part of the FlexConfig object.
- Determine what types of variables you will need, and create any policy objects that you will require. You cannot create objects for variables while editing a FlexConfig object.
- Ensure that your commands do not conflict in any way with the VPN or access control configuration on the devices.
- If there is more than one set of commands for an interface, only the last set of commands is deployed. Therefore, we recommend you not use beginning and ending commands to configure interfaces. For an example of configuring interfaces, see the ISIS_Interface_Configuration predefined FlexConfig object.

Procedure

Step 1 Choose **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object**.

Step 2 Do one of the following:

- Click **Add FlexConfig Object** to create a new object.
- Click **Edit** ( ) to edit an existing object.
- Click **View** ( ) to see the contents of a predefined object.
- If you want to edit a predefined object, click **Clone** ( ) to create a new object with the same contents.

Step 3 Enter a **Name** and optionally, a description for the object.

Step 4 In the object body area, enter the commands and instructions to produce the required configuration.

The object content is a sequence of scripting commands and configuration commands that generate a valid ASA software command sequence. The Firewall Threat Defense device uses ASA software commands to configure some features. For more information on scripting and configuration commands, see:

- Template Scripts, on page 2671

- CLI Commands in FlexConfig Objects, on page 2668

You can use variables to supply information that can be known only at runtime, or which can differ from device to device. You simply type in processing variables, but you must use the **Insert** menu to add variables that are associated with policy objects or system variables, or which are secret keys. For a complete discussion of variables, see FlexConfig Variables, on page 2672.

- To insert system variables, choose **Insert** > **Insert System Variable** > ***Variable Name***. For a detailed explanation of these variables, see FlexConfig System Variables, on page 2677.
- To insert policy object variables, choose **Insert** > **Insert Policy Object** > ***Object Type***, selecting the appropriate type of object. Then, give the variable a name (which can be the same name as the associated policy object), select the object to associate with the variable, and click **Save**. For a detailed explanation of these types, see FlexConfig Policy Object Variables, on page 2676. For more detail on the procedure, see Add a Policy Object Variable to a FlexConfig Object, on page 2692.
- To insert secret key variables, choose **Insert** > **Secret Key** and define the variable name and value. For more detail on the procedure, see Configure Secret Keys, on page 2693.

Note
You must use the **Insert** menu to create a new policy object or system variable. However, for subsequent uses of that variable, you must type it in, $ included. This is also true for system variables: the first time you use it, add it from the **Insert** menu. Then, type it out for subsequent uses. If you use the **Insert** menu more than once for a system variable, the system variable is added to the Variables list multiple times, and the FlexConfig will not validate, meaning you cannot save your changes. For processing variables (those not associated with a policy object or system variable), simply type in the variable. If you are adding a secret key, always use the **Insert** menu. Secret key variables do not show up in the Variables list.

Step 5 Choose the deployment frequency and type.

- **Deployment**—Whether to deploy the commands in the object **Once** or **Everytime**. The only way to choose the right option is to test the results of deployment.

  Start by selecting **Everytime**. Then, after you attach the object to a FlexConfig policy, deploy the configuration. After a successful deployment, come back to the FlexConfig policy and preview the configuration for one of the assigned devices as described in Preview the FlexConfig Policy, on page 2697. If the section labeled `###CLI generated from managed features ###` contains commands that clear or negate the commands in the object, and the `###Flex-config Appended CLI ###` section contains the commands to reconfigure the feature, you know that **Everytime** is the right option.

  Even if you do not see negate commands, make some minor change to the device configuration, then run another deployment. If the deployment completes successfully, you can check the deployment transcript (see Verify the Deployed Configuration, on page 2698). If you see that the commands were issued again (even when they were already configured) without error, then you can keep **Everytime**.

  Change to **Once** only if the system does not first negate the commands in the object before issuing them again, or if the deployment results in errors that are specific to the command. In some cases, the system does not allow you to issue a command that is already configured, but this is the exception.

  Some additional tips:

  - If the FlexConfig object points to system-managed objects such as network or ACL objects, choose **Everytime**. Otherwise, updates to the objects might not get deployed.
  - Use **Once** if the only thing you do in the object is to clear a configuration. Then, remove the object from the FlexConfig policy after the next deployment.

- **Type**—Select one of the following:
  - **Append**—(The default.) Commands in the object are put at the end of the configurations generated from the Firewall Management Center policies. You must use Append if you use policy object variables, which point to objects generated from managed objects. If commands generated for other policies overlap with those specified in the object, you should select this option so your commands are not overwritten. This is the safest option.
  - **Prepend**—Commands in the object are put at the beginning of the configurations generated from the Firewall Management Center policies. You would typically use prepend for commands that clear or negate a configuration.

Step 6 (Optional.) Click **Validate** ( ) above the object body to check the integrity of the script.

The object is always validated when you click **Save**. You cannot save an invalid object.

Step 7 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Add a Policy Object Variable to a FlexConfig Object

You can insert variables into a FlexConfig policy object that are associated with other types of policy object. When the FlexConfig is deployed to a device, these variables resolve to the names or content of the associated object.

Use the following procedure for the first use of a policy object variable in a FlexConfig object. If you need to refer to the object again, type in the variable (including the $ sign). To understand how to use these variables, see How to Process Variables, on page 2672.

Before you begin

For information on editing a FlexConfig Object, see Configure FlexConfig Objects, on page 2690.

Procedure

Step 1 While editing a FlexConfig Policy Object, choose **Insert** > **Insert Policy Object** > ***Object Type***, selecting the appropriate type of object.

Step 2 Enter a name for the variable, and optionally, a description.

The name must be unique within the context of the FlexConfig object. It cannot include spaces. You are allowed to use the exact same name as the object associated with the variable.

Step 3 Select the object to associate with the variable and click **Add** to move it to the **Selected Object** list.

You can associate a variable with a single object only.

Note

For text objects, you can select any of the predefined objects as needed. However, many of these objects have no default values. You must update the objects to add the required values either directly or as overrides for the device to which you will deploy the FlexConfig object. Trying to deploy a FlexConfig without updating these objects typically results in deployment errors.

Step 4 Click **Save**.

The variable appears in the Variables list at the bottom of the FlexConfig object editor.

Configure Secret Keys

A secret key is any single-string variable whose content you want to mask, such as passwords. The system provides special treatment for these variables to help you prevent the dissemination of sensitive information.

Secret key variables do not show up in the Variables list in the FlexConfig object.

Use the following procedure to create, insert, and otherwise manage secret key variables in a FlexConfig object. Unlike other types of variables, you can use the **Insert** command every time you need to insert a given secret key variable. With respect to processing, these variables behave like single-value text object variables; see Single Value Variables, on page 2672.

Note Any data defined in a secret key variable is masked from users except when previewing a FlexConfig policy. In addition, if you export a FlexConfig policy, the content of any secret key variable is erased. When you import the policy, you will need to manually edit each secret key variable to enter the data.

Before you begin

For information on editing a FlexConfig Object, see Configure FlexConfig Objects, on page 2690.

Procedure

Step 1 While editing a FlexConfig Policy Object, choose **Insert** > **Secret Key**.

Step 2 In the **Insert Secret Key** dialog box, do any of the following:

- To create a new key, click **Add Secret Key**, then fill in the following information and click **Add**.
  - **Secret Key Name**—The name of the variable. This name appears in the FlexConfig object prefixed with @.
  - **Password**, **Confirm Password**—The secret string, which is masked with asterisks as you type.
- To insert a secret key variable in the FlexConfig object, select the check box for the variable.
- To edit the value of a secret key variable, click **Edit** () for the variable. Make your changes and click **Add**.
- To delete a secret key variable, click **Delete** () for the variable.

Step 3 Click **Save**.

Configure FlexConfig Text Objects

Use text objects in FlexConfig objects as the target of policy object variables. You can use variables to supply information that can be known only at runtime, or which can differ from device to device. During deployment, variables that point to text objects are replaced by the content of the text object.

Text objects contain free-form strings, which can be keywords, interface names, numbers, IP addresses, and so forth. The content depends on how you will use the information within a FlexConfig script.

Before creating or editing a text object, determine exactly what content you will need. This includes how you intend to process the object, which will help you decide between creating a single string or multiple string object. Read the following topics:

- FlexConfig Variables, on page 2672
- How to Process Variables, on page 2672

Procedure

Step 1 Choose **Objects** > **Object Management** > **FlexConfig** > **Text Object**.

Step 2 Do one of the following:

- Click **Add Text Object** to create a new object.
- Click **Edit** (✎) to edit an existing object. You are allowed to edit the predefined text objects, which is required if you intended to use the predefined FlexConfig objects.

Step 3 Enter a **Name** and optionally, a description for the object.

Step 4 (New objects only.) Choose a **Variable Type** from the drop-down list:

- **Single**—If the object should contain a single text string.
- **Multiple**—If the object should contain a list of text strings.

You cannot change the variable type after you save the object.

Step 5 If the variable type is **Multiple**, use the up and down arrows to specify a **Count**.

Rows are added or removed from the object as you change the number.

Step 6 Add content to the object.

You can either click in the text box next to a variable number and type in a value, or you can set up device overrides for each device that will be assigned a FlexConfig object that uses the text object. You can also do both, in which case the values configured in the base object act as default values in cases where an override does not exist for a given device.

When editing predefined objects, it is a good practice to use device overrides, so that the system defaults remain in place for other users who might need to use the object in different FlexConfig policies. The approach you take depends on the requirements of your organization.

Tip
Some predefined objects require multiple values where each value serves a specific purpose. Read the description text carefully to determine the expected values in the object. In some cases, the instructions specify that you must use overrides instead of changing the base values. In the case of enableInspectProtocolList, you are prevented from entering protocols whose inspection is incompatible with Snort inspection.

If you decide to use device overrides, do the following.

a) Check the check box of **Allow Overrides**.
b) Expand the Overrides area (if necessary) and click **Add**.

   If an override already exists for the device, click edit for the override to change it.

c) On **Targets** in the Add Object Override dialog box, select the device for which you are defining values and click **Add** to move it to the Selected Devices list.
d) Click **Override**, adjust the **Count** as needed, then click in the variable fields and type in the values for the device.
e) Click **Add**.

Step 7 Click **Save**.

What to do next

- If an active policy references your object, deploy configuration changes; see Deploy Configuration Changes, on page 259.

Configure the FlexConfig Policy

A FlexConfig policy contains two ordered lists of FlexConfig objects, one prepended list and one appended list. For an explanation of prepend/append, see Configure FlexConfig Objects, on page 2690.

FlexConfig policies are shared policies that you can assign to multiple devices.

Procedure

Step 1 Choose **Devices** > **FlexConfig**.

Step 2 Do one of the following:

- Click **New Policy** to create a new FlexConfig Policy. You are prompted to enter a name. Optionally, select devices in the Available Devices list and click **Add to Policy** to assign devices. Click **Save**.
- Click **Edit** ( ) to edit an existing Policy. You can change the name or description by clicking them in edit mode.
- Click **Copy** ( ) to create a new policy with the same contents. You are prompted for a name. Device assignments are not retained for the copy.
- Click delete to remove a policy you no longer need.

Step 3 Select the FlexConfig objects required for the policy from the **Available FlexConfig** list and click > to add them to the policy.

Objects are automatically added to the prepended or appended list based on the deployment type specified in the FlexConfig object.

To remove a selected object, click **Delete** ( ) next to an object.

Step 4 For each selected object, click **View** ( ) next to the object to identify the variables used in the object.

Except for system variables, which start with SYS, you need to ensure that the objects associated with the variables are not empty. A blank or brackets with nothing between them, [ ], indicate an empty object. You will need to edit these objects before deploying the policy.

Note
If you use object overrides, those values will not show up in this view. Thus, an empty default value does not necessarily mean that you have not updated the object with the required values. Previewing the configuration will show whether the variables resolve correctly for a given device. See Preview the FlexConfig Policy, on page 2697.

Step 5 Click **Save**.

What to do next

- Set target devices for the policy; see Set Target Devices for a FlexConfig Policy, on page 2696.
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Set Target Devices for a FlexConfig Policy

When you create a FlexConfig policy, you can select the devices that use the policy. You can subsequently change device assignments for the policy as described below.

Note Normally, when you unassign a policy from a device, the system automatically removes the associated configuration upon the next deployment. However, because FlexConfig objects are scripts for deploying customized commands, simply unassigning a FlexConfig policy from a device does not remove the commands that were configuring by the FlexConfig objects. If your intention is to remove FlexConfig-generated commands from a device's configuration, see Remove Features Configured Using FlexConfig, on page 2699.

Procedure

Step 1 Choose **Devices** > **FlexConfig** and edit a FlexConfig policy.

Step 2 Click **Policy Assignments**.

Step 3 On **Targeted Devices**, build your target list:

- Add—Choose one or more **Available Devices**, then click **Add to Policy** or drag and drop into the list of **Selected Devices**. You can assign the policy to devices, high availability pairs, and clustered devices.
- Delete—Click **Delete** ( ) next to a single device, or select multiple devices, right-click, then choose **Delete Selection**.

Step 4 Click **OK** to save your selection.

Step 5 Click **Save** to save the FlexConfig policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Preview the FlexConfig Policy

Preview a FlexConfig policy to see how the FlexConfig objects get translated into CLI commands. The preview shows the commands that will be generated for a selected device from the scripts and variables used in the FlexConfig objects. The variables are resolved based on the configuration for the device, so you get a clear idea of what will be deployed.

Use the preview to look for potential problems in the FlexConfig objects. Correct the objects until the preview shows the expected results.

You must preview the configuration separately for each device, because the variables can resolve differently based on the device configuration.

Procedure

Step 1 Choose **Devices** > **FlexConfig** and edit a FlexConfig policy.

Step 2 If there are any pending changes, click **Save**.

The preview shows results only for those FlexConfig objects that were in the most recently saved version of the policy. You must save the policy to see a preview of newly-added objects.

Step 3 Click **Preview Config**.

Step 4 Choose a device from the **Select Device** drop-down list.

The system retrieves information from the device and configured policies, and determines what CLI commands will be generated on the next deployment to the device. You can select the output and use Ctrl+C to copy it to the clipboard, where you can paste it into a text file for further analysis.

The preview includes the following sections:

- Flex-config Prepended CLI—These are the commands generated by FlexConfigs that are prepended to the configuration.
- CLI generated from managed features—These are commands generated for policies configured in the Firewall Management Center. Commands are generated for new or changed policies since the last successful deployment to the device. These commands do not represent all commands needed to implement the assigned policies. No commands in this section are generated from FlexConfig objects.

- Flex-config Appended CLI—These are the commands generated by FlexConfigs that are appended to the configuration.

Step 5 Click **Close** to close the preview dialog.

Verify the Deployed Configuration

After you deploy a FlexConfig policy to a device, verify that the deployment was successful and that the resulting configuration is what you expected. Also, verify that the device is performing as expected.

Procedure

Step 1 To verify that deployment was successful:

a) Click **Notifications** in the menu bar, which is unnamed between **Deploy** and **System**.

The icon looks like one of the following, and it might include a number if there are errors:

- **Indicates No Warnings** — Indicates no warnings or errors are present on the system.
- Indicates **One or More Warnings** — Indicates one or more warnings and no errors are present on the system.
- **Indicates One or More Errors** — Indicates one or more errors and any number of warnings are present on the system.

b) On **Deployments**, verify that the deployment was successful.

c) To see more detailed information, especially for failed deployments, click **Show History**.

d) Select the deployment job in the list of jobs in the left column.

Jobs are listed in reverse chronological order, with the most recent job at the top of the list.

e) Click download in the **Transcript** column for the device in the right column.

The deployment transcript includes commands sent to the device, and any responses returned from the device. These response can be informative messages or error messages. For failed deployments, look for messages that indicate errors with the commands that you sent through FlexConfig. These errors can help you correct the script in the FlexConfig object that is trying to configure the commands.

Note
There is no distinction made in the transcript between commands sent for managed features and those generated from FlexConfig policies.

For example, the following sequence shows that Firewall Management Center sent commands to configure GigabitEthernet0/0 with the logical name outside. The device responded that it automatically set the security level to 0. Firewall Threat Defense does not use the security level for anything. Messages relevant to FlexConfig are in the CLI Apply section of the transcript.

```
========= CLI APPLY =========

FMC >> interface GigabitEthernet0/0
FMC >>  nameif outside
```

```
FTDv 192.168.0.152 >> [info] : INFO: Security level for "outside" set to 0 by default.
```

Step 2 Verify that the deployed configuration includes the expected commands.

You can do this by making an SSH connection to the device's management IP address. Use the **show running-config** command to view the configuration.

Alternatively, use the CLI tool within Secure Firewall Management Center.

a) Choose **System** (⚙) > **Health** > **Monitor** and click the name of the device.

You might need to click the open/close arrow in the **Count** column in the Status table to see any devices.

b) Click **Advanced Troubleshooting**.
c) Click **Threat Defense CLI**.
d) Select **show** as the command, and type **running-config** as the parameter.
e) Click **Execute**.

The running configuration appears in the text box. You can select the configuration and press Ctrl+C, then paste it into a text file for later analysis.

Step 3 Verify that the device is performing as expected.

Use the **show** commands related to the feature to see detailed information and statistics. For example, if you enabled additional protocol inspections, the **show service-policy** command provides this information. The exact commands to use are feature-dependent and should be mentioned in the ASA configuration guide and command reference you used to learn how to configure the feature.

If commands that show statistics indicate that numbers are not changing (for example, hit counts, connection counts, and so forth), the configuration might be valid but not meaningful. If you know that traffic is going through the device that should show up in statistics, look for what is missing in your configuration. For example, NAT or access rules might be dropping or changing traffic before a feature can act on it.

You can use the **show** commands from an SSH session or through the Firewall Management Center CLI tool.

However, if the **show** command that you need to use is not available directly within the Firewall Threat Defense CLI, you will need make an SSH connection to the device to use the commands. From the CLI, enter the following command sequence to enter Privileged EXEC mode within the diagnostic CLI. From there, you should be able to enter these otherwise unsupported **show** commands.

```
> system support diagnostic-cli
Attaching to Diagnostic CLI ... Press 'Ctrl+a then d' to detach.
Type help or '?' for a list of available commands.
firepower> enable
Password: <press enter, do not enter a password>
firepower#
```

Remove Features Configured Using FlexConfig

If you decide you need to remove a set of configuration commands you configured using FlexConfig, you might need to manually remove that configuration. Unassigning the FlexConfig policy from a device might not remove all of the configuration.

To manually remove the configuration, you create new FlexConfig objects to clear or negate the configuration commands.

Before you begin

To determine if you need to manually remove some or all of the configuration generated by an object:

1. Examine the configuration preview, as described in Preview the FlexConfig Policy, on page 2697. If the `###CLI generated from managed features ###` section contains the clear or negate commands to remove all of the commands in the FlexConfig object, then you can simply remove the object from the FlexConfig policy, save, and redeploy.
2. Remove the object from the FlexConfig policy, save the change, then preview the configuration again. If the `###CLI generated from managed features ###` section still does not include the required clear or negate commands, you must follow this procedure to manually remove the configuration.

Procedure

Step 1 Choose **Objects** > **Object Management** and create the FlexConfig Objects to clear or negate the configuration commands.

If a feature has a **clear** command that can remove all configuration settings, then use that command. For example, the predefined ISIS_Unconfigure_All object contains a single command that removes all ISIS-related configuration commands:

```
clear configure router isis
```

If there is not a **clear** command for the feature, you need to use the **no** form of each command you want to remove. For example, the predefined Sysopt_basic_negate object removes the commands configured through the predefined Sysopt_basic object.

```
no sysopt traffic detailed-statistics
```

```
no sysopt connection timewait
```

You would typically configure a FlexConfig object that removes configurations as a prepended, deploy once object.

Step 2 Choose **Devices** > **FlexConfig** and create a new FlexConfig policy or edit the existing policy.

If you want to preserve the FlexConfig policy that deploys the configuration commands, create a new policy specifically for negating the commands, and assign the devices to the policy. Then, add the new FlexConfig objects to the policy.

If you want to completely remove the FlexConfig configuration objects from all devices, you can simply delete those commands from the existing FlexConfig policy and replace them with the objects that negate the configuration.

Step 3 Click **Save** to save the FlexConfig policy.

Step 4 Click **Preview Config** and verify that the clear and negation commands are generating correctly.

Step 5 Choose **Deploy** > **Deployment** in the menu bar, select the device, and click **Deploy**.

Wait for deployment to complete.

Step 6 Verify that the commands were removed.

View the running configuration on the device to confirm that the commands are removed. For more detailed information, see Verify the Deployed Configuration, on page 2698.

Step 7 While editing the FlexConfig policy, click **Policy Assignments** and remove the device. Optionally, remove the FlexConfig Objects from the policy.

Assuming that the FlexConfig policy simply removes the unwanted configuration commands, there is no need to keep the policy assigned to the device after the removal is complete.

However, if the FlexConfig policy retains options that you still want configured on the device, remove the negation objects from the policy. They are no longer needed.

Convert from FlexConfig to Managed Feature

Each software release adds managed features to the product, that is, features that you configure directly through policies that are controlled outside of FlexConfig. This can deprecate FlexConfig commands that you are currently using; your configurations are not automatically converted. After the upgrade, you cannot assign or create FlexConfig objects using the newly deprecated commands. After upgrading software, examine your FlexConfig policies and objects.

When a feature you configured using FlexConfig starts to be supported as a managed feature, you must convert from using FlexConfig to using the managed feature. In most cases, your existing FlexConfig configurations continue to work post-upgrade and you can still deploy. However, in some cases, using deprecated commands can cause deployment issues. Configuring a feature in both the GUI and FlexConfig is not supported.

Note Use the migration tool instead of this procedure if the tool supports the feature configuration you are migrating.

Procedure

Step 1 Remove the FlexConfig, as explained in Remove Features Configured Using FlexConfig, on page 2699.

Step 2 Configure the settings in the newly supported managed feature.

The release notes have a list of new features for the release.

Examples for FlexConfig

Following are some examples of using FlexConfig.

How to Configure Precision Time Protocol (ISA 3000)

The Precision Time Protocol (PTP) is a time-synchronization protocol developed to synchronize the clocks of various devices in a packet-based network. These device clocks are generally of varying precision and stability. The protocol is designed specifically for industrial, networked measurement and control systems, and is optimal for use in distributed systems because it requires minimal bandwidth and little processing overhead.

A PTP system is a distributed, networked system consisting of a combination of PTP and non-PTP devices. PTP devices include ordinary clocks, boundary clocks and transparent clocks. Non-PTP devices include network switches, routers and other infrastructure devices.

You can configure the Firewall Threat Defense device to be a transparent clock. The Firewall Threat Defense device does not synchronize its clock with the PTP clocks. The Firewall Threat Defense device will use the PTP default profile, as defined on the PTP clocks.

When you configure the PTP devices, you define a domain number for the devices that are meant to function together. Thus, you can configure multiple PTP domains, and then configure each non-PTP device to use the PTP clocks for one specific domain.

Before you begin

Determine the domain number configured on the PTP clocks that the device should use. This example assumes the PTP domain number is 10. Also, determine the interfaces through which the system can reach the PTP clocks in the domain.

Following are guidelines for configuring PTP:

- This feature is only available on the Cisco ISA 3000 appliance.
- Cisco PTP supports multicast PTP messages only.
- PTP is available only for IPv4 networks, not for IPv6 networks.
- PTP configuration is supported on physical Ethernet data interfaces, whether stand-alone or bridge group members. It is not supported on the management interface, subinterfaces, EtherChannels, Bridge Virtual Interfaces (BVI), or any other virtual interfaces.
- PTP flows on VLAN subinterfaces are supported, assuming the appropriate PTP configuration is present on the parent interface.
- You must ensure that PTP packets are allowed to flow through the device. PTP traffic is identified by UDP destination ports 319 and 320, and destination IP address 224.0.1.129, so any access control rule that allows this traffic should work.
- In Routed firewall mode, you must enable Multicast routing for PTP multicast groups. In addition, if an interface on which you enable PTP is **not** in a bridge group, you must configure the interface to join the IGMP multicast group 224.0.1.129. If the physical interface is a bridge group member, you do not configure it to join the IGMP multicast group.

Procedure

Step 1 (Routed mode only.) Enable Multicast routing, and configure the IGMP group for the interfaces.

In Routed mode, you must enable Multicast routing. In addition, for stand-alone physical interfaces, that is, those that are not bridge group members, you must also configure the interface to join the 224.0.1.129 IGMP group. You cannot configure bridge group members to join an IGMP group, but PTP configuration on bridge group members will work without the IGMP join.

Perform this procedure for each device on which you will configure PTP.

Note
Write down the hardware names of each PTP-clock-facing interface on each device, for example, GigabitEthernet1/1.

a) Choose **Devices** > **Device Management**, and edit the device.
b) Click **Routing**.
c) Choose **Multicast Routing** > **IGMP**.
d) Check the **Enable Multicast Routing** check box.
e) Click **Join Group**.
f) Click **Add**, and in the **Add IGMP Join Group Parameters** dialog box, configure the following options and click **OK**.

- **Interface**—Select the PTP-clock-facing stand-alone interface.
- **Join Group**—Click **+** to add a new network object. Create a Host object with the address 224.0.1.129. When configuring additional interfaces, simply select this object. (See Creating Network Objects, on page 1403.)

Repeat this step for each PTP-clock-facing stand-alone interface on the device.

g) Click **Save** on the Routing page.

Step 2 Create the FlexConfig object to enable PTP globally and on the interface.

The following procedure assumes that the PTP-clock-facing interface is the same on every device you are configuring. If you have used different interfaces on different devices, you need to create separate objects for each distinct combination. For example, if you use GigabitEthernet1/1 on devices A and B, GigabitEthernet1/2 on devices C and D, and both GigabitEthernet1/1 and 1/2 on devices E and F, you need 3 separate FlexConfig objects, and subsequently, 3 separate FlexConfig policies (explained in the next step).

a) Choose **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object**.
b) Click **Add FlexConfig Object**, configure the following properties, and click **Save**.

- **Name**—The object name. For example, Enable_PTP.
- **Deployment**—Select **Everytime**. You want this configuration to be sent in every deployment to ensure it remains configured.
- **Type**—Keep the default, **Append**. The commands are sent to the device after the commands for directly-supported features. This ensures that any other changes you make to interface configuration are configured before these commands.
- **Object body**—In the object body, type the commands needed to configure PTP globally and on each PTP-clock-facing interface. For example, the commands needed for the global configuration for PTP domain 10 and the interface configuration on GigabitEthernet1/1 are:

```
ptp mode e2etransparent
ptp domain 10
```

```
interface gigabitethernet1/1
 ptp enable
```

The object body should look similar to the following:

Step 3 Create the FlexConfig policy and assign it to the devices.

If you created multiple FlexConfig objects for different combinations of PTP-clock-facing interfaces, you need to create separate FlexConfig policies for each object, and assign those policies to the correct devices based on the interfaces you need to configure. Repeat the following procedure for each group of devices.

a) Choose **Devices** > **FlexConfig**.

b) Either click **New Policy**, or if an existing FlexConfig policy should be assigned to (or is already assigned to) the target devices, simply edit that policy.

When creating a new policy, assign the target devices to the policy in the dialog box where you name the policy.

c) Select the PTP FlexConfig object in the **User Defined** folder in the table of contents and click > to add it to the policy.

The object should be added to the **Selected Appended FlexConfigs** list.

d) Click **Save**.

e) If you have not yet assigned all the targeted devices to the policy, click the **Policy Assignments** link below Save and make the assignments now.

f) Click **Preview Config**, and in the Preview dialog box, select one of the assigned devices.

The system generates a preview of the configuration CLI that will be sent to the device. Verify that the commands generated from the PTP FlexConfig object look correct. These will be shown at the end of the

preview. Note that you will also see commands generated from other changes you have made to managed features. For the PTP commands, you should see something similar to the following:

```
###Flex-config Appended CLI ###
ptp mode e2etransparent
ptp domain 10
interface gigabitethernet1/1
 ptp enable
```

Step 4 Deploy your changes.

Because you assigned a FlexConfig policy to the devices, you will always get a deployment warning, which is meant to caution you about the use of FlexConfig. Click **Proceed** to continue with the deployment.

After the deployment completes, you can check the deployment history and view the transcript for the deployment. This is especially valuable if the deployment fails. See Verify the Deployed Configuration, on page 2698.

Step 5 Verify the PTP configuration on each device.

From an SSH or Console session into each device, verify the PTP settings:

```
> show ptp clock
 PTP CLOCK INFO
  PTP Device Type: End to End Transparent Clock
  Operation mode: One Step
  Clock Identity: 34:62:88:FF:FE:1:73:81
  Clock Domain: 10
  Number of PTP ports: 4
> show ptp port
 PTP PORT DATASET: GigabitEthernet1/1
  Port identity: Clock Identity: 34:62:88:FF:FE:1:73:81
  Port identity: Port Number: 1
  PTP version: 2
  Port state: Enabled

 PTP PORT DATASET: GigabitEthernet1/2
  Port identity: Clock Identity: 34:62:88:FF:FE:1:73:81
  Port identity: Port Number: 2
  PTP version: 2
  Port state: Disabled

 PTP PORT DATASET: GigabitEthernet1/3
  Port identity: Clock Identity: 34:62:88:FF:FE:1:73:81
  Port identity: Port Number: 3
  PTP version: 2
  Port state: Disabled

 PTP PORT DATASET: GigabitEthernet1/4
  Port identity: Clock Identity: 34:62:88:FF:FE:1:73:81
  Port identity: Port Number: 4
  PTP version: 2
  Port state: Disabled
```

How to Configure Automatic Hardware Bypass for Power Failure (ISA 3000)

You can enable hardware bypass so that traffic continues to flow between an interface pair during a power outage. Supported interface pairs are copper interfaces GigabitEthernet 1/1 and 1/2; and GigabitEthernet 1/3 and 1/4. If you have a fiber Ethernet model, only the copper Ethernet pair (GigabitEthernet 1/1 and 1/2) supports hardware bypass.

When hardware bypass is active, traffic passes between these interface pairs at layer 1. The Firewall Threat Defense CLI will see the interfaces as being down. No firewall functions are in place, so make sure you understand the risks of allowing traffic to pass through the device.

In CLI Console or an SSH session, use the **show hardware-bypass** command to monitor the operational status.

Before you begin

For hardware bypass to work:

- You must place the interface pairs in the same bridge group.
- You must attach the interfaces to access ports on the switch. Do not attach them to trunk ports.

We recommend that you disable TCP sequence number randomization globally using the Threat Defense Service Policy attached to the access control policy assigned to the device. By default, the ISA 3000 rewrites the initial sequence number (ISN) of TCP connections passing through it to a random number. When hardware bypass is activated, the ISA 3000 is no longer in the data path and does not translate the sequence numbers. The receiving client receives an unexpected sequence number and drops the connection, so the TCP session needs to be re-established. Even with TCP sequence number randomization disabled, some TCP connections will have to be re-established because of the link that is temporarily down during the switchover.

Procedure

Step 1 Create the FlexConfig object to enable automatic bypass.

a) Choose **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Object**.

b) Click **Add FlexConfig Object**, configure the following properties, and click **Save**.

- **Name**—The object name. For example, Enable_HW-Bypass.
- **Deployment**—Select **Everytime**. You want this configuration to be sent in every deployment to ensure it remains configured.
- **Type**—Keep the default, **Append**. The commands are sent to the device after the commands for directly-supported features.
- **Object body**—In the object body, type the commands needed to enable automatic hardware bypass. For example, the commands needed for both possible interface pairs:

```
hardware-bypass GigabitEthernet 1/1-1/2
hardware-bypass GigabitEthernet 1/3-1/4
```

The object body should look similar to the following:

Step 2 Create the FlexConfig policy and assign it to the devices.

a) Choose **Devices** > **FlexConfig**.

b) Either click **New Policy**, or if an existing FlexConfig policy should be assigned to (or is already assigned to) the target devices, simply edit that policy.

When creating a new policy, assign the target devices to the policy in the dialog box where you name the policy.

c) Select the hardware bypass FlexConfig object in the **User Defined** folder in the table of contents and click > to add it to the policy.

The object should be added to the **Selected Appended FlexConfigs** list.

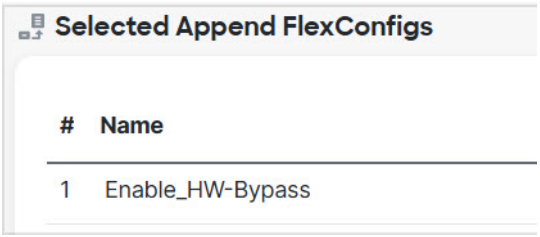

d) Click **Save**.

e) If you have not yet assigned all the targeted devices to the policy, click the **Policy Assignments** link below Save and make the assignments now.

f) Click **Preview Config**, and in the Preview dialog box, select one of the assigned devices.

The system generates a preview of the configuration CLI that will be sent to the device. Verify that the commands generated from the hardware bypass FlexConfig object look correct. These will be shown at the end of the preview. Note that you will also see commands generated from other changes you have made to managed features. For the hardware bypass commands, you should see something similar to the following:

```
###Flex-config Appended CLI ###
hardware-bypass GigabitEthernet 1/1-1/2
hardware-bypass GigabitEthernet 1/3-1/4
```

Step 3 Deploy your changes.

Because you assigned a FlexConfig policy to the devices, you will always get a deployment warning, which is meant to caution you about the use of FlexConfig. Click **Proceed** to continue with the deployment.

After the deployment completes, you can check the deployment history and view the transcript for the deployment. This is especially valuable if the deployment fails. See Verify the Deployed Configuration, on page 2698.

What to do next

If you want to manually invoke hardware bypass or manually turn it off, you need to create two FlexConfig objects:

- One that manually starts bypass, which would contain one or both of the following commands, depending on whether you want to invoke bypass for both pairs:

```
hardware-bypass manual GigabitEthernet 1/1-1/2
hardware-bypass manual GigabitEthernet 1/3-1/4
```

- One that manually turns off bypass, which would contain one or both of the following commands:

```
no hardware-bypass manual GigabitEthernet 1/1-1/2
no hardware-bypass manual GigabitEthernet 1/3-1/4
```

You would then need to add one or the other object to the FlexConfig policy, and deploy changes, to either turn bypass on or off. You would also need to immediately remove the object from the FlexConfig policy after deployment. If you manually invoke bypass, you would then need to repeat the process to turn it off again. Thus, using this manual method requires frequent and careful editing of the FlexConfig policy and additional deployments.

Migrating FlexConfig Policies

Attention This section on migrating FlexConfig policies is applicable only for migrating ECMP, VXLAN, and EIGRP policies.

The ECMP, VXLAN, and EIGRP policies were configured using the FlexConfig objects and policies in earlier versions of the Firewall Management Center. You can now directly configure these policies in the management center UI. When you are upgrading the management center from earlier versions, the FlexConfig configuration is retained. However, to manage the policies from the UI, you must redo the configuration in the corresponding **Device (Edit)** > **Routing** page and remove the configuration from FlexConfig. To automate creation of the policies in the UI, Firewall Management Center provides an option to migrate the policies from FlexConfig to the UI. However, it does not remove the migrated policies from the FlexConfig. For post-migration procedure, refer to Step 7, on page 2709.

Before you begin

- Ensure that the deployed FlexConfig policy is up-to-date and not out-of-date. The migration option will be available only if the policy is up-to-date on at least 1 device. Migration does not take place for those devices that have out-of-date policies.
- If the policy is configured in both FlexConfig and the management center:

- Migration will not be initiated if the policy is already configured at **Device (Edit)** > **Routing**.
- During deployment, the management center displays an error message. Example EIGRP migration error message—*EIGRP is configured through FlexConfig object and also under Device Listing ->Routing EIGRP for the device. Maintain the EIGRP configuration in either Routing EIGRP or FlexConfig.*

- If the network objects used in the policy exist in the management center, during migration, they are reused. During migration, when a network object that is matching the IP configuration is not available, a new network object is created as *bb* appended with a timestamp and integer, like ***bb_<timestamp>_<integer>***. For more than one such network object, the integer variable in the name would be incremented by one.

Procedure

Step 1 Choose **Devices** > **FlexConfig**, click **Edit** (✎) against the FlexConfig policy that you want to migrate.

Step 2 Click **Migrate Config**.

Note
Once migration starts, both **Migrate Config** and FlexConfig **Edit** options are not available.

The **Migrate Config** option is not available, when:

- There are no applicable FlexConfig CLIs to be migrated.
- The FlexConfig policy is not associated with any FlexConfig objects.
- There are no devices that are associated with the FlexConfig policy.

Step 3 In the **Migrate Flex Configuration** dialog box, select the device to which you wish to migrate the configuration, and then click **Ok**.

The migration progress is displayed as task notification. When the migration is complete, click the *View Details* link and download the migration report (PDF format).

Step 4 To view the policy changes, choose **System** (⚙) > **Monitoring** > **Audit** and click on the *Flex Config Migration* message.

Step 5 To view the FlexConfig migration report, choose **System** (⚙) > **Monitoring** > **Audit** and click on the *Flex Config Migration* message. To view the complete migration report, click the **Report** icon.

Step 6 Verify the migrated configuration settings in the corresponding **Device (Edit)** > **Routing**. page.

Step 7 To remove the specific policy config from FlexConfig for the device, in the management center, do the following:

a) Identify the migrated FlexConfig policy for the device.
b) Use the copy option and create duplicate of the FlexConfig policy.
c) Remove the corresponding CLI objects from the duplicated FlexConfig policy.
d) Associate the device to the duplicated FlexConfig policy.

Step 8 Save and deploy the configuration.

History for FlexConfig

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Quick recovery after data plane failure for the Firepower 1000/2100 and Firepower 4100/9300.	7.4.1	7.4.1	If the data plane process crashes, the system now reloads only the data plane process instead of rebooting the device. Along with the data plane process reload, Snort and a few other processes also get reloaded. However, if the data plane process crashes during bootup, the device follows the normal reload/reboot sequence, which helps avoid a reload process loop from occurring. This feature is enabled by default for both new and upgraded devices. New/modified CLI commands: **data-plane quick-reload**, **no data-plane quick-reload**, **show data-plane quick-reload status** Supported platforms: Firepower 1000/2100, Firepower 4100/9300 Platform restrictions: Not supported in multi-instance mode. See: Cisco Secure Firewall Threat Defense Command Reference and Cisco Secure Firewall ASA Series Command Reference.
Migration tool support.	7.3.0	Any	Support to migrate Flex configured ECMP, VXLAN, and EIGRP policies to management center was introduced. New/modified screens: **Devices** > **FlexConfig** > **Migrate FlexConfig**
Removal of BFD configuration in FlexConfig.	7.3.0	Any	Support to configure BFD policies directly in the management center user interface was introduced. Hence, FlexConfig support to configure BFD policies was removed.
Removal of priority-queue.	7.2.5	7.2.5	Support to configure priority-queue in threat defense was removed.
Removal of EIGRP configuration in FlexConfig.	7.2.0	Any	Support to configure EIGRP directly in the management center user interface was introduced. Hence, FlexConfig support to configure EIGRP policies was removed.
Removal of PBR configuration.	7.1.0	7.1.0	Support to configure the PBR directly in the FMC user interface was introduced. Hence, FlexConfig support to configure PBR for FTD 7.1 and higher was removed. New/modified commands: **policy-route route-map***routemap-object-name*.
Removal of ECMP zone creation support in FlexConfig.	7.1.0	Any	Support to configure the ECMP zones directly in the FMC user interface was introduced. Hence, FlexConfig support to configure ECMP zones was removed.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Precision Time Protocol (PTP) configuration for ISA 3000 devices.	6.5.0	Any	You can use FlexConfig to configure the Precision Time Protocol (PTP) on ISA 3000 devices. PTP is a time-synchronization protocol developed to synchronize the clocks of various devices in a packet-based network. The protocol is designed specifically for industrial, networked measurement and control systems. We now allow you to include the **ptp** (interface mode) command, and the global commands **ptp mode e2etransparent** and **ptp domain**, in FlexConfig objects. New/modified commands: **show ptp**.
Deprecated FlexConfig objects.	6.3.0	Any	Several features that in previous releases you configured using FlexConfig are now directly supported in FMC. You need to remove these FlexConfig objects if you are using them, and convert your configuration to use the new objects. Following are the deprecated FlexConfig objects and text objects. • **Default_DNS_Configure**, including the defaultDNSNameServerList and defaultDNSParameters text objects. Now, please configure DNS for the data interfaces using the Platform Settings policy. • **TCP_Embryonic_Conn_Limit**, and the tcp_conn_misc and tcp_conn_limit text objects. Configure these features in the FTD Service Policy, which you can find on the Advanced tab of the access control policy assigned to the device. • **TCP_Embryonic_Conn_Timeout**, and the tcp_conn_misc and tcp_conn_timeout text objects. Configure these features in the FTD Service Policy.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
FlexConfig updates.	6.2.1	Any	As per the Government Certification requirements, all sensitive information like password, shared keys in system-provided or user-defined FlexConfig object should be masked using secret key variables. After you update the FMC to Version 6.2.1+, all sensitive information in FlexConfig Objects are converted to secret key variable format. In addition, the following new FlexConfig templates are added: • **Default_DNS_Configure** template allows you to the default DNS group, which is used to resolve hostnames for commands or features that resolve names through the data interfaces. • **TCP Embryonic connection limit and timeout configuration** template allows you to configure embryonic connection limits/timeout CLIs to protect from SYN Flood DoSAttack. • **Turn on threat detection configure and clear** templates allow you to configure threat detection statistics for attacks intercepted by TCP Intercept. • **IPV6 router header inspection** template allows you to configure of IPV6 inspection header for selectively allow/block certain headers with different types (e.g. allowing RH Type 2,mobile). • **DHCPv6 prefix delegation** template allows you to configure one outside (PD client) and one inside interface (recipient of delegated prefix) for IPv6 prefix delegation.
FlexConfig.	6.2.0	Any	The FlexConfig feature allows you use the FMC to deploy ASA CLI template-based functionality to FTD devices. This feature allows you to enable some of the most valuable ASA functions that are not currently available on FTD devices. This functionality is structured as templates and objects that work together in a policy. The default templates are officially supported by Cisco TAC. New/modified screens: • **Devices** > **FlexConfig** • **Objects** > **Object Management** > **FlexConfig** > **FlexConfig Objects** • **Objects** > **Object Management** > **FlexConfig** > **Text Object**

PART XV

Advanced Network Analysis in Snort 3

- Advanced Access Control Settings for Network Analysis and Intrusion Policies, on page 2715
- Transport and Network Layer Preprocessors, on page 2723
- Adaptive Profiles, on page 2727

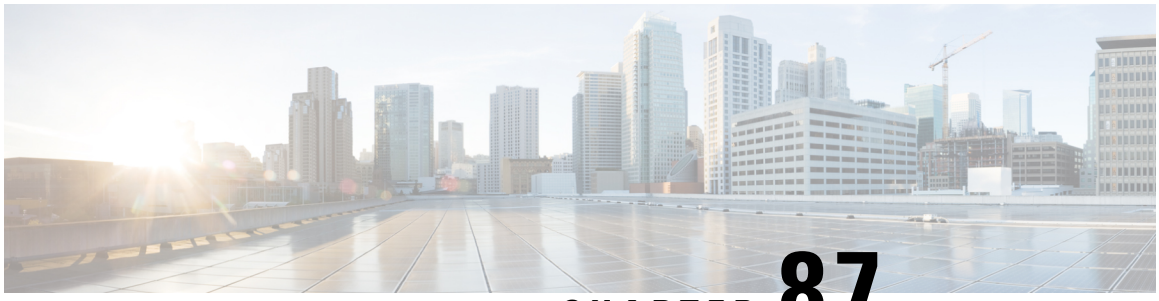

CHAPTER 87

Advanced Access Control Settings for Network Analysis and Intrusion Policies

The following topics describe how to configure advanced settings for network analysis and intrusion policies:

- About Advanced Access Control Settings for Network Analysis and Intrusion Policies, on page 2715
- Requirements and Prerequisites for Advanced Access Control Settings for Network Analysis and Intrusion Policies, on page 2715
- Inspection of Packets That Pass Before Traffic Is Identified, on page 2716
- Advanced Settings for Network Analysis Policies, on page 2717

About Advanced Access Control Settings for Network Analysis and Intrusion Policies

Many of the advanced settings in an access control policy govern intrusion detection and prevention configurations that require specific expertise to configure. Advanced settings typically require little or no modification and are not common to every deployment.

Note Snort 2 is not supported on threat defense Version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the Firewall Management Center guide that matches your Firewall Threat Defense version.

Requirements and Prerequisites for Advanced Access Control Settings for Network Analysis and Intrusion Policies

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Inspection of Packets That Pass Before Traffic Is Identified

For some features, including URL filtering, application detection, rate limiting, and Intelligent Application Bypass, a few packets must pass in order for the connection to be established, and to enable the system to identify the traffic and determine which access control rule (if any) will handle that traffic.

You must explicitly configure your access control policy to inspect these packets, prevent them from reaching their destination, and generate any events.

As soon as the system identifies the access control rule or default action that should handle the connection, the remaining packets in the connection are handled and inspected accordingly.

Best Practices for Handling Packets That Pass Before Traffic Identification

- The default action specified for an access control policy is NOT applied to these packets.
- Instead, use the following guidelines to choose a value for the **Intrusion Policy used before Access Control rule is determined** setting in the Advanced settings of the access control policy.
  - You can choose a system-created or custom intrusion policy. For example, you can choose **Balanced Security and Connectivity**.
  - For performance reasons, unless you have good reason to do otherwise, this setting should match the default action set for your access control policy.
  - If your system does not perform intrusion inspection (for example, in a discovery-only deployment), select **No Rules Active**. The system will not inspect these initial packets, and they will be allowed to pass.
  - By default, this setting uses the default variable set. Ensure that this is suitable for your purposes. For information, see Variable Set, on page 1446.
  - The network analysis policy associated with the first matching network analysis rule preprocesses traffic for the policy you select. If there are no network analysis rules, or none match, the default network analysis policy is used.

Specify a Policy to Handle Packets That Pass Before Traffic Identification

Note This setting is sometimes referred to as the *default intrusion policy*. (This is distinct from the default action for an access control policy.)

Before you begin

Review best practices for these settings. See Best Practices for Handling Packets That Pass Before Traffic Identification, on page 2716.

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** () next to the **Network Analysis** and **Intrusion Policies** section.

If **View** () appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 Select an intrusion policy from the **Intrusion Policy used before Access Control rule is determined** drop-down list.

If you choose a user-created policy, you can click **Edit** () to edit the policy in a new window. You cannot edit system-provided policies.

Step 3 Optionally, select a different variable set from the **Intrusion Policy Variable Set** drop-down list. You can also select **Edit** () next to the variable set to create and edit variable sets. If you do not change the variable set, the system uses a default set.

Step 4 Click **OK**.

Step 5 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Related Topics

Variable Set, on page 1446

Advanced Settings for Network Analysis Policies

Network analysis policies govern how traffic is decoded and preprocessed so that it can be further evaluated, especially for anomalous traffic that might signal an intrusion attempt. This traffic preprocessing occurs after Security Intelligence matching and traffic decryption, but before intrusion policies inspect packets in detail. By default, the system-provided Balanced Security and Connectivity network analysis policy is the default network analysis policy.

Tip The system-provided Balanced Security and Connectivity network analysis policy and the Balanced Security and Connectivity intrusion policy work together and can both be updated in intrusion rule updates. However, the network analysis policy governs mostly preprocessing options, whereas the intrusion policy governs mostly intrusion rules.

A simple way to tune preprocessing is to create and use a custom network analysis policy as the default. For advanced users with complex deployments, you can create multiple network analysis policies, each tailored to preprocess traffic differently. Then, you can configure the system to use those policies to govern the preprocessing of traffic using different security zones, networks, or VLANs.

To accomplish this, you add custom *network analysis rules* to your access control policy. A network analysis rule is simply a set of configurations and conditions that specifies how you preprocess traffic that matches those qualifications. You create and edit network analysis rules in the advanced options in an existing access control policy. Each rule belongs to only one policy.

Each rule has:

- a set of rule conditions that identifies the specific traffic you want to preprocess
- an associated network analysis policy that you want to use to preprocess traffic that meets all the rules' conditions

When it is time for the system to preprocess traffic, it matches packets to network analysis rules in top-down order by rule number. Traffic that does not match any network analysis rules is preprocessed by the default network analysis policy.

Setting the Default Network Analysis Policy

You can choose a system- or user-created policy.

Note If you disable a preprocessor but the system needs to evaluate preprocessed packets against an enabled intrusion or preprocessor rule, the system automatically enables and uses the preprocessor although it remains disabled in the network analysis policy web interface. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task. Because preprocessing and intrusion inspection are so closely related, you **must** be careful that you allow the network analysis and intrusion policies examining a single packet to complement each other.

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** ( ) next to the Network Analysis and Intrusion Policies section.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 From the **Default Network Analysis Policy** drop-down list, select a default network analysis policy.

If you choose a user-created policy, you can click **Edit** (✎) to edit the policy in a new window. You cannot edit system-provided policies.

Step 3 Click **OK**.

Step 4 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Related Topics

Limitations of Custom Policies

Network Analysis Rules

Within your access control policy's advanced settings, you can use network analysis rules to tailor preprocessing configurations to network traffic.

Network analysis rules are numbered, starting at 1. When it is time for the system to preprocess traffic, it matches packets to network analysis rules in top-down order by ascending rule number, and preprocesses traffic according to the first rule where all the rule's conditions match.

You can add zone, network, and VLAN tag conditions to a rule. If you do not configure a particular condition for a rule, the system does not match traffic based on that criterion. For example, a rule with a network condition but no zone condition evaluates traffic based on its source or destination IP address, regardless of its ingress or egress interface. Traffic that does not match any network analysis rules is preprocessed by the default network analysis policy.

Network Analysis Policy Rule Conditions

Rule conditions enable you to fine-tune your network analysis policy to target the users and networks you want to control. See one of the following sections for more information.

Related Topics

Security Zone Rule Conditions, on page 1905
Network Rule Conditions, on page 939
VLAN Tags Rule Conditions, on page 1787

Security Zone Rule Conditions

Security zones segment your network to help you manage, classify, and decrypt traffic flow by grouping interfaces across multiple devices.

Security zones control or decrypt traffic by its source and destination security zones. If you add both source and destination zones to a zone condition, matching traffic must originate from an interface in one of the source zones and leave through an interface in one of the destination zones.

Just as all interfaces in a zone must be of the same type (all inline, passive, switched, or routed), all zones used in a zone condition must be of the same type. Because devices deployed passively do not transmit traffic, you cannot use a zone with passive interfaces as a destination zone.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Tip Constraining rules by zone is one of the best ways to improve system performance. If a rule does not apply to traffic through any of device's interfaces, that rule does not affect that device's performance.

Security Zone Conditions and Multitenancy

In a multidomain deployment, a zone created in an ancestor domain can contain interfaces that reside on devices in different domains. When you configure a zone condition in an descendant domain, your configurations apply to only the interfaces you can see.

Network Rule Conditions

Networks control or decrypt traffic by its source and destination IP address, using inner headers. Tunnel rules, which use outer headers, have tunnel endpoint conditions instead of network conditions.

You can use predefined objects to build network conditions, or manually specify individual IP addresses or address blocks.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Note You *cannot* use FDQN network objects in identity rules.

VLAN Tags Rule Conditions

Note VLAN tags in access rules only apply to inline sets. Access rules with VLAN tags do not match traffic on firewall interfaces.

VLAN rule conditions control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic. The system uses the innermost VLAN tag to filter VLAN traffic, with the exception of the prefilter policy, which uses the outermost VLAN tag in its rules.

Note the following Q-in-Q support:

- Firewall Threat Defense on Firepower 4100/9300—Does not support Q-in-Q (supports only one VLAN tag).
- Firewall Threat Defense on all other models:
  - Inline sets and passive interfaces—Supports Q-in-Q, up to 2 VLAN tags.
  - Firewall interfaces—Does not support Q-in-Q (supports only one VLAN tag).

You can use predefined objects to build VLAN conditions, or manually enter any VLAN tag from 1 to 4094. Use a hyphen to specify a range of VLAN tags.

In a cluster, if you encounter problems with VLAN matching, edit the access control policy advanced options, Transport/Network Preprocessor Settings, and select the **Ignore the VLAN header when tracking connections** option.

Configuring Network Analysis Rules

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** (✎) next to the Network Analysis and Intrusion Policies section.

If **View** (👁) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings.If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Tip
Click **Network Analysis Policy List** to view and edit existing custom network analysis policies.

Step 2 Next to **Network Analysis Rules**, click the statement that indicates how many custom rules you have.

Step 3 Click **Add Rule**.

Step 4 Configure the rule's conditions by clicking the conditions you want to add.

Step 5 Click **Network Analysis** and choose the **Network Analysis Policy** you want to use to preprocess the traffic matching this rule.

Click **Edit** (✎) to edit a custom policy in a new window. You cannot edit system-provided policies.

Step 6 Click **Add**.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Managing Network Analysis Rules

A network analysis rule is simply a set of configurations and conditions that specifies how you preprocess traffic that matches those qualifications. You create and edit network analysis rules in the advanced options in an existing access control policy. Each rule belongs to only one policy.

Procedure

Step 1 In the access control policy editor, click **Advanced**, then click **Edit** (✎) next to the Intrusion and Network Analysis Policies section.

If **View** (👁) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings.If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 2 Next to **Network Analysis Rules**, click the statement that indicates how many custom rules you have.

Step 3 Edit your custom rules. You have the following options:

- To edit a rule's conditions, or change the network analysis policy invoked by the rule, click **Edit** ( ) next to the rule.
- To change a rule's order of evaluation, click and drag the rule to the correct location. To select multiple rules, use the Shift and Ctrl keys.
- To delete a rule, click **Delete** ( ) next to the rule.

Tip
Right-clicking a rule displays a context menu that allows you to cut, copy, paste, edit, delete, and add new network analysis rules.

Step 4 Click **OK**.

Step 5 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

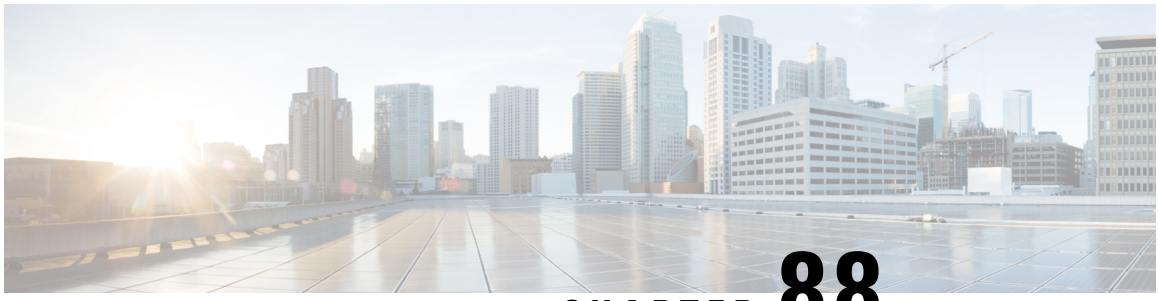

CHAPTER 88

Transport and Network Layer Preprocessors

The following topics explain transport and network layer preprocessors and how to configure them:

- Introduction to Transport and Network Layer Preprocessors, on page 2723
- License Requirements for Transport and Network Layer Preprocessors, on page 2723
- Requirements and Prerequisites for Transport and Network Layer Preprocessors, on page 2724
- Advanced Transport/Network Preprocessor Settings, on page 2724

Introduction to Transport and Network Layer Preprocessors

Transport and network layer preprocessors detect attacks that exploit IP fragmentation, checksum validation, and TCP and UDP session preprocessing. Before packets are sent to preprocessors, the packet decoder converts packet headers and payloads into a format that can be easily used by the preprocessors and the intrusion rules engine and detects various anomalous behaviors in packet headers. After packet decoding and before sending packets to other preprocessors, the inline normalization preprocessor normalizes traffic for inline deployments.

When an intrusion rule or rule argument requires a disabled preprocessor, the system automatically uses it with its current configuration even though it remains disabled in the network analysis policy's web interface.

Note Snort 2 is not supported on threat defense Version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the Firewall Management Center guide that matches your Firewall Threat Defense version.

License Requirements for Transport and Network Layer Preprocessors

Threat Defense License

IPS

Requirements and Prerequisites for Transport and Network Layer Preprocessors

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Intrusion Admin

Advanced Transport/Network Preprocessor Settings

Advanced transport and network preprocessor settings apply globally to all networks, zones, and VLANs where you deploy your access control policy. You configure these advanced settings in an access control policy rather than in a network analysis policy.

Ignored VLAN Headers

Different VLAN tags in traffic traveling in different directions for the same connection can affect traffic reassembly and rule processing. For example, in the following graphic traffic for the same connection could be transmitted over VLAN A and received over VLAN B.

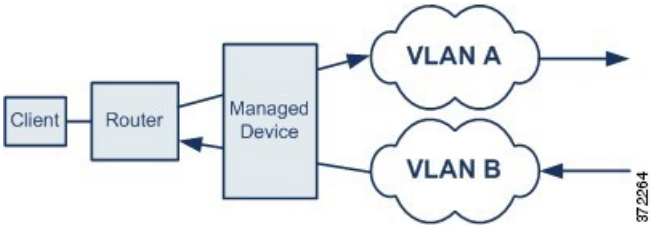

You can configure the system to ignore the VLAN header so packets can be correctly processed for your deployment.

Active Responses in Intrusion Drop Rules

A drop rule is an intrusion or preprocessor rule whose rule state is set to Drop and Generate Events. In an inline deployment, the system responds to TCP or UDP drop rules by dropping the triggering packet and blocking the session where the packet originated.

Tip Because UDP data streams are not typically thought of in terms of *sessions*, the stream preprocessor uses the source and destination IP address fields in the encapsulating IP datagram header and the port fields in the UDP header to determine the direction of flow and identify a UDP session.

You can configure the system to initiate one or more *active responses* to more precisely and specifically close a TCP connection or UDP session when an offending packet triggers a TCP or UDP drop rule. You can use active responses in inline, including routed and transparent, deployments. Active responses are not suited or supported for passive deployments.

To configure active responses:

- Create or modify a TCP or UDP (**resp** keyword only) intrusion rule.
- Add the **react** or **resp** keyword to the intrusion rule; see xActive Response Keywords.
- Optionally, for a TCP connection, specify the maximum number of additional active responses to send and the number of seconds to wait between active responses; see **Maximum Active Responses** and **Minimum Response Seconds** in Advanced Transport/Network Preprocessor Options, on page 2725.

Active responses close the session when matching traffic triggers a drop rule, as follows:

- **TCP**—drops the triggering packet and inserts a TCP Reset (RST) packet in both the client and server traffic.
- **UDP**—sends an ICMP unreachable packet to each end of the session.

Advanced Transport/Network Preprocessor Options

Ignore the VLAN header when tracking connections

Specifies whether to ignore or include VLAN headers when identifying traffic, as follows:

- When this option is selected, the system ignores VLAN headers. Use this setting for deployed devices that might detect different VLAN tags for the same connection in traffic traveling in different directions
- When this option is disabled, the system includes VLAN headers. Use this setting for deployed devices that will not detect different VLAN tags for the same connection traffic traveling in different directions.

Maximum Active Responses

Specifies a maximum number of active responses per TCP connection. When additional traffic occurs on a connection where an active response has been initiated, and the traffic occurs more than **Minimum Response Seconds** after a previous active response, the system sends another active response unless the specified maximum has been reached. A setting of 0 disables additional active responses triggered by **resp** or **react** rules. See Active Responses in Intrusion Drop Rules, on page 2724 and Active Response Keywords.

Note that a triggered **resp** or **react** rule initiates an active response regardless of the configuration of this option.

Minimum Response Seconds

Until **Maximum Active Responses** occur, specifies the number of seconds to wait before any additional traffic on a connection where the system has initiated an active response results in a subsequent active response.

Troubleshooting Options: Session Termination Logging Threshold

Caution Do not modify Session Termination Logging Threshold unless instructed to do so by Support.

Support might ask you during a troubleshooting call to configure your system to log a message when an individual connection exceeds the specified threshold. Changing the setting for this option will affect performance and should be done only with Support guidance.

This option specifies for the number of bytes that result in a logged message when the session terminates and the specified number was exceeded.

Note The upper limit of 1GB is also restricted by the amount of memory on the managed device allocated for stream processing.

Related Topics

Active Response Keywords

Configuring Advanced Transport/Network Preprocessor Settings

You must be an Admin, Access Admin, or Network Admin to perform this task.

Procedure

Step 1 In the access control policy editor, click **Edit** (✎) on the policy you want to modify.

Step 2 Click **More** > **Advanced Settings**, and then click **Edit** (✎) next to the **Transport/Network Preprocessor Settings** section.

Step 3 Except for the troubleshooting option **Session Termination Logging Threshold**, modify the options described in Advanced Transport/Network Preprocessor Options, on page 2725.

Caution
Do not modify **Session Termination Logging Threshold** unless instructed to do so by Support.

Step 4 Click **OK**.

What to do next

- Optionally, further configure the policy as described in Editing an Access Control Policy, on page 1749.
- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

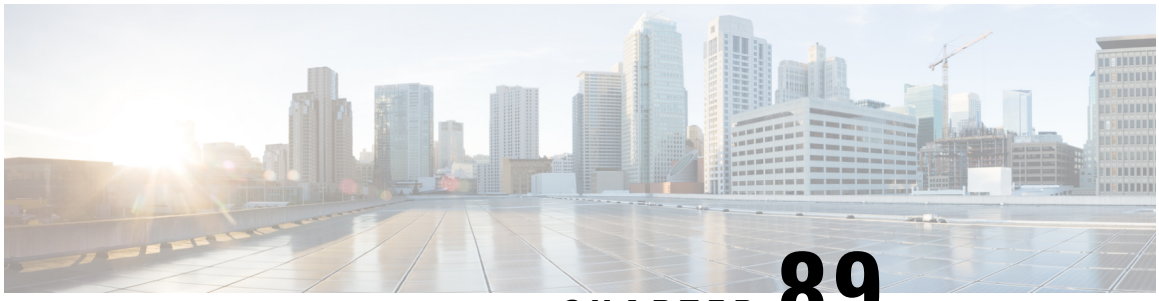

CHAPTER 89

Adaptive Profiles

The following topics describe how to configure adaptive profiles:

- About Adaptive Profiles, on page 2727
- License Requirements for Adaptive Profiles, on page 2728
- Requirements and Prerequisites for Adaptive Profiles, on page 2728
- Adaptive Profile Updates, on page 2728
- Adaptive Profile Updates and Cisco Recommended Rules, on page 2729
- Adaptive Profile Options, on page 2729
- Configuring Adaptive Profiles, on page 2730

About Adaptive Profiles

Adaptive profiles must be enabled in order to:

- Perform application and file control, including malware protection (AMP), and to allow intrusion rules to use service metadata.

Caution Adaptive profiling **must** be enabled (its default state) as described in Configuring Adaptive Profiles, on page 2730 for access control rules to perform application and file control, including malware protection (AMP), and for intrusion rules to use service metadata.

- For passive deployments, enable adaptive profile updates to defragment and reassemble IP traffic according to the destination hosts' operating systems.

Note For inline deployments Cisco recommends that, instead of enabling adaptive profile updates, you configure the inline normalization preprocessor with the **Normalize TCP Payload** option enabled.

Note Snort 2 is not supported on threat defense Version 7.7. For information on Snort 2 features that are supported in versions earlier than 7.7, refer to the Firewall Management Center guide that matches your Firewall Threat Defense version.

License Requirements for Adaptive Profiles

Threat Defense License

IPS

Requirements and Prerequisites for Adaptive Profiles

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Adaptive Profile Updates

Typically, the system uses the static settings in your network analysis policy to preprocess and analyze traffic. With adaptive profile updates, the system can adapt processing behavior using host information either detected by network discovery or imported from a third party.

Profile updates, like the target-based profiles you can configure manually in a network analysis policy, help to defragment IP packets and reassemble streams in the same way as the operating system on the target host. The intrusion rules engine then analyzes the data in the same format as that used by the destination host.

Manually configured target-based profiles apply either the default operating system profile you select, or profiles you bind to specific hosts. Profile updates, however, switch to the appropriate operating system profile based on the operating system in the host profile for the target host.

Consider a scenario where you configure profile updates for the 10.6.0.0/16 subnet and set the default IP Defragmentation target-based policy to Linux. The Firewall Management Center where you configure the settings has a network map that includes the 10.6.0.0/16 subnet.

- When the system detects traffic from Host A, which is not in the 10.6.0.0/16 subnet, it uses the Linux target-based policy to reassemble IP fragments.
- When the system detects traffic from Host B, which is in the 10.6.0.0/16 subnet, it retrieves Host B's operating system data from the network map. The system uses a profile based on that operating system to defragment the traffic destined for Host B.

Adaptive Profile Updates and Cisco Recommended Rules

The adaptive profile updates feature is an advanced setting in an access control policy that applies globally to all intrusion policies invoked by that access control policy. The Cisco recommended rules feature applies to the individual intrusion policy where you configure it.

Like Cisco recommended rules, profile updates compare metadata in a rule to host information to determine whether a rule should apply for a particular host. However, while Cisco recommended rules provide recommendations for enabling or disabling rules using that information, profile updates use the information to apply specific rules to specific traffic.

Cisco recommended rules require your interaction to implement suggested changes to rule states. Profile updates, on the other hand, do not modify intrusion policies. Treatment of rules based on profile updates happens on a packet-by-packet basis.

Additionally, Cisco recommended rules can result in enabling disabled rules. Profile updates, in contrast, only affect the application of rules that are already enabled in intrusion policies. Profile updates never change the rule state.

You can use profile updates and Cisco recommended rules in combination. Profile updates use the rule state for a rule when your intrusion policy is deployed to determine whether to include it as a candidate for applying, and your choices to accept or decline recommendations are reflected in that rule state. You can use both features to ensure that you have enabled or disabled the most appropriate rules for each network you monitor, and then to apply enabled rules most efficiently for specific traffic.

Related Topics

About Cisco Recommended Rules

Adaptive Profile Options

Enable

Enabling this option is required for:

- access control rules to perform application and file control, including malware protection (AMP)
- intrusion rules to use service metadata

This option is enabled by default.

Note To enable Adaptive Profiles in Snort 3, both **Enable** and **Enable Profile Updates** options must be selected.

Enable Profile Updates

In passive deployments, enable profile updates to defragment and reassemble IP traffic according to a profile of the operating system used by the hosts in your network map

For Snort 3, this must be enabled if Adaptive Profiles is enabled.

Adaptive Profiles - Attribute Update Interval

When profile updates are enabled, you can control how frequently in minutes network map data is synced from the Firewall Management Center to its managed devices. The system uses the data to determine what profiles should be used when processing traffic. Increasing the value for this option can improve performance in a large network.

Adaptive Profiles - Networks

Optionally, when profile updates are enabled, you can improve performance by constraining profile updates to a comma-separated list of IP addresses, address blocks, and network variables. If you use a network variable, the system uses the variable's value in the variable set linked to the default intrusion policy for your access control policy. For example, you could enter: **`192.168.1.101, 192.168.4.0/24, $HOME_NET`**. IPv4 and IPv6 are supported.

The default value (**`0.0.0.0/0`**) applies adaptive profile updates to all networks.

Related Topics

Inspection of Packets That Pass Before Traffic Is Identified, on page 2716
Variable Set, on page 1446

Configuring Adaptive Profiles

In a passive deployment, Cisco recommends that you configure adaptive profile updates. In an inline deployment, configure the inline normalization preprocessor with the **Normalize TCP Payload** option enabled.

Caution Adaptive profiling **must** be enabled (its default state) as described in this procedure for access control rules to perform application or file control, including AMP, and for intrusion rules to use service metadata.

Before you begin

The access control policy must have a network discovery policy that is enabled to do host/service discovery, or host data must be imported from a third-party source.

Procedure

Step 1 In the access control policy editor, click **Edit** (✎) on the policy you want to modify.

Step 2 Click **More** > **Advanced Settings**, and then click **Edit** (✎) next to the **Detection Enhancement Settings** section.

If **View** ( ) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings.If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Set adaptive profile options as described in Adaptive Profile Options, on page 2729.

Step 4 Click **OK**.

Step 5 Click **Save** to save the policy.

What to do next

- Deploy configuration changes; see Deploy Configuration Changes, on page 259.

Related Topics

Snort Restart Scenarios, on page 252

PART XVI

Threat Intelligence Director

- Secure Firewall Threat Intelligence Director, on page 2735

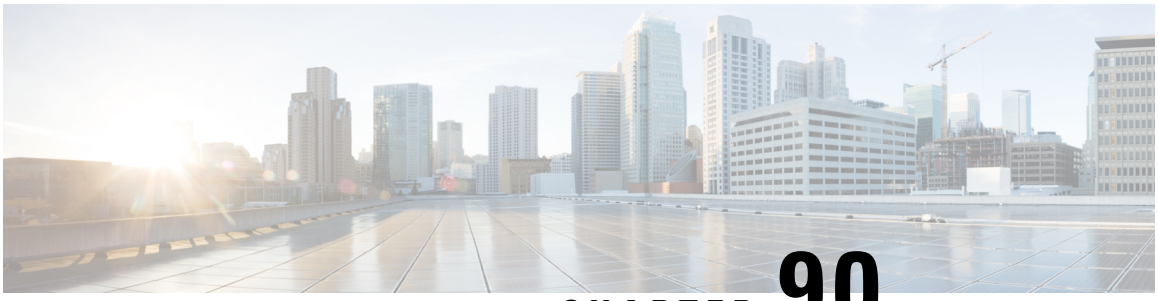

CHAPTER 90

Secure Firewall Threat Intelligence Director

The topics in this chapter describe how to configure and use Threat Intelligence Director.

- Secure Firewall Threat Intelligence Director Overview, on page 2735
- Requirements and Prerequisites for Threat Intelligence Director, on page 2738
- How To Set Up Threat Intelligence Director, on page 2740
- Analyze Threat Intelligence Director Incident and Observation Data, on page 2749
- View and Change Threat Intelligence Director Configurations, on page 2761
- Troubleshoot Threat Intelligence Director, on page 2776
- History for Threat Intelligence Director, on page 2778

Secure Firewall Threat Intelligence Director Overview

The Secure Firewall Threat Intelligence Director operationalizes threat intelligence data, helping you aggregate intelligence data, configure defensive actions, and analyze threats in your environment. This feature is intended to supplement other Firepower functionality, offering an additional line of defense against threats.

When configured on your hosting platform, Threat Intelligence Director ingests data from threat intelligence *sources* and publishes the data to all configured managed devices (*elements*.) For more information about the hosting platforms and elements supported in this release, see Platform, Element, and License Requirements, on page 2738.

Sources contain *indicators*, which contain *observables*. An indicator conveys all of the characteristics associated with a threat, and individual observables represent individual characteristics (e.g. a SHA-256 value) associated with the threat. *Simple indicators* contain a single observable, and *complex indicators* contain two or more observables.

Observables and the AND/OR operators between them form an indicator's *pattern*, as illustrated in the following examples.

Figure 450: Example: Indicator Patterns

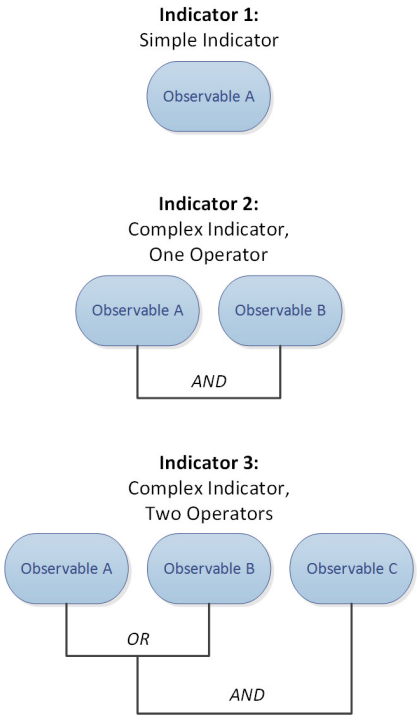

After the observables are published to the elements, the elements monitor traffic and report *observations* to the Firewall Management Center when the system identifies observables in traffic.

The Firewall Management Center collects observations from all elements, evaluates the observations against Threat Intelligence Director indicators, and generates or updates *incidents* associated with the observable's parent indicator(s).

An incident is *fully realized* when an indicator's pattern is fulfilled. An incident is *partially realized* if traffic matches one or more observables in the indicator but not the entire pattern. For more information, see Observation and Incident Generation, on page 2749.

The following diagram shows data flow in a sample system configuration.

Figure 451: Firewall Management Center Data Flow

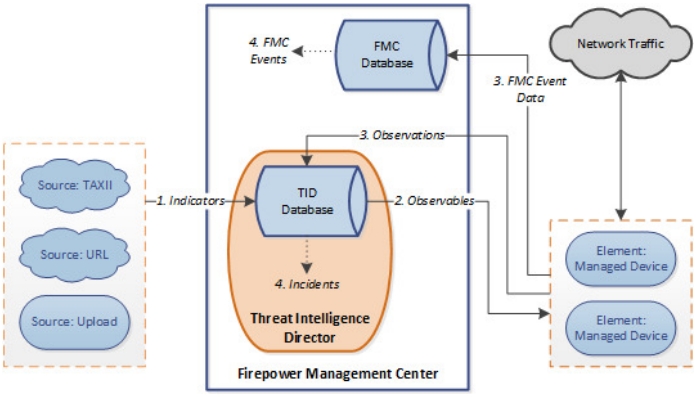

When a Threat Intelligence Director incident is fully or partially realized, the system takes the configured *action* (monitor, block, partially block, or no action). For details, see Factors That Affect the Action Taken, on page 2757.

Threat Intelligence Director and Security Intelligence

As part of your access control policy, Security Intelligence uses reputation intelligence to quickly block connections to or from IP addresses, URLs, and domains. Security Intelligence uniquely provides access to industry-leading threat intelligence from Talos Intelligence Group. For more information on Security Intelligence, see About Security Intelligence, on page 1885.

Threat Intelligence Director enhances the system's ability to block connections based on security intelligence from third-party sources as follows:

- **Threat Intelligence Director supports additional traffic filtering criteria**—Security Intelligence allows you to filter traffic based on IP address, URL, and (if DNS policy is enabled) domain name. Threat Intelligence Director also supports filtering by these criteria and adds support for filtering on SHA-256 hash values.
- **Threat Intelligence Director supports additional intelligence ingestion methods**—With both Security Intelligence and Threat Intelligence Director, you can import threat intelligence into the system by either manually uploading flat files or configuring the system to retrieve flat files from a third-party host. Threat Intelligence Director provides increased flexibility in managing those flat files. In addition, Threat Intelligence Director can retrieve and ingest intelligence provided in Structured Threat Information eXpression (STIX™) format.
- **Threat Intelligence Director provides granular control of filtering actions**—With Security Intelligence, you can specify filtering criteria by network, URL, or DNS object. Security Intelligence objects, especially list and feeds, can contain multiple IP addresses, URLs, or DNS domain names, but you can only block or not block based on entire objects, not based on individual components of an object. With Threat Intelligence Director, you can configure filtering actions for individual criteria (that is, simple indicators or individual observables).
- **Threat Intelligence Director configuration changes do not require redeployment**—After you modify Security Intelligence settings in the access control policy, you must redeploy the changed configuration to managed devices. With Threat Intelligence Director, after initial deployment of the access control policy to the managed devices, you can configure sources, indicators, and observables without redeploying, and the system automatically publishes new Threat Intelligence Director data to the elements.

For information about what the system does when either Security Intelligence or Threat Intelligence Director could handle a particular incident, see Threat Intelligence Director-Firewall Management Center Action Prioritization, on page 2757.

Performance Impact of Threat Intelligence Director

Secure Firewall Management Center

In some cases, you may notice the following:

- The system may experience minor performance issues while ingesting particularly large STIX sources, and ingestion may take longer than expected to finish.

- The system may take up to 15 minutes to publish new or modified Threat Intelligence Director data down to elements.

Managed Device

There is no exceptional performance impact. Threat Intelligence Director impacts performance identically to the Secure Firewall Management Center Security Intelligence feature.

Requirements and Prerequisites for Threat Intelligence Director

Model Support

Any

Supported Domains

Any

User Roles

Admin

Threat Intelligence Director User

Additional Requirements

The following topics explain additional requirements for using Threat Intelligence Director.

Platform, Element, and License Requirements

Hosting Platforms

You can host Threat Intelligence Director on physical and virtual Secure Firewall Management Centers:

- running Version 6.2.2 or later.
- configured with a minimum of 15 GB of memory.
- configured with REST API access enabled. See *Enabling REST API Access* in the Cisco Secure Firewall Management Center Administration Guide.

Elements

You can use any Secure Firewall Management Center-managed device as a Threat Intelligence Director element if the device is running Version 6.2.2 or later.

Licensing

To configure the file policies for SHA-256 observable publishing, you need the following licensed devices:

- For smart licensed devices:

- IPS License - For IPv4, IPv6, URL, and DNS detection and observables
- Malware Defense License - For SHA-256 detection and observables

For more information, see Configure Policies to Support Threat Intelligence Director, on page 2741 and the *Licenses* chapter in the Cisco Secure Firewall Management Center Administration Guide.

Source Requirements

Source Type Requirements:

STIX

Files must be STIX Version 1.0, 1.1, 1.1.1, or 1.2 and adhere to the guidelines in the STIX documentation: http://stixproject.github.io/documentation/suggested-practices/.

STIX files can include complex indicators.

The maximum size for a STIX file is 40MB when configured via URL download or file upload. If you have STIX files larger than this, we recommend using a TAXII server.

Flat File

Files must be ASCII text files with one observable value per line.

Flat files include only simple indicators (one observable per indicator.)

Flat files can be up to 500 MB.

Threat Intelligence Director does **not** support:

- Delimiter characters separating observable values (e.g. `observable,` is invalid).
- Enclosing characters around observable values (e.g. `"observable"` is invalid).

Each file should contain only one type of content:

- `SHA-256`—SHA-256 hash values.
- `Domain`—domain names as defined in RFC 1035.
- `URL`—URLs as defined in RFC 1738.

Note Threat Intelligence Director normalizes any URLs that contain port, protocol, or authentication information, and uses the normalized version when detecting indicators. For example, Threat Intelligence Director normalizes any of the following URLs:

```
http://example.com/index.htm
http://example.com:8080/index.htm
example.com:8080/index.htm
example.com/index.htm
```

as:

```
example.com/index.htm
```

Or, for example, Threat Intelligence Director normalizes the following URL:

```
http://abc@example.com:8080/index.htm
```

as

```
abc@example.com/index.htm/
```

- `IPv4`—IPv4 addresses as defined in RFC 791.

  Threat Intelligence Director does not accept CIDR blocks.

- `IPv6`—IPv6 addresses as defined in RFC 4291.

  Threat Intelligence Director does not accept prefix lengths.

Source Content Limitations

The system ingests, and matches on, only the first 1000 characters of a URL observable.

How To Set Up Threat Intelligence Director

Note If you encounter an issue during Threat Intelligence Director configuration or operation, see Troubleshoot Threat Intelligence Director, on page 2776.

Procedure

Step 1 Ensure that your installation meets the requirements for running Threat Intelligence Director.

See Platform, Element, and License Requirements, on page 2738

Step 2 For each managed device, configure the policies required to support Threat Intelligence Director and deploy those policies to the devices.

See Configure Policies to Support Threat Intelligence Director, on page 2741.

You can configure elements before or after you ingest intelligence data sources.

Step 3 Configure the intelligence sources that you want Threat Intelligence Director to ingest.

See Source Requirements, on page 2739 and the topics under Options for Ingesting Data Sources, on page 2742.

Step 4 Publish data to the elements if you have not yet done so. See Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level, on page 2774.

What to do next

- Include Threat Intelligence Director in your regularly scheduled backups. See About Backing Up and Restoring Threat Intelligence Director Data, on page 2748.

  If your Secure Firewall Management Center deployment is a high availability configuration, see also *Firewall Management Center High Availability Disaster Recovery* in the Cisco Secure Firewall Management Center Administration Guide.
- (Optional) Grant administrative access to Threat Intelligence Director functionality as desired. See User Roles with Threat Intelligence Director Access, on page 2748 and the *Users* chapter in the Cisco Secure Firewall Management Center Administration Guide.
- As needed during operation, fine-tune your configuration. For example, add observables that produce false-positive incidents to the Do Not Block list. See View and Change Threat Intelligence Director Configurations, on page 2761.

Configure Policies to Support Threat Intelligence Director

You must configure access control policies to publish Threat Intelligence Director data from the Firewall Management Center to your managed devices (elements). In addition, we recommend that you configure your access control policies to maximize observation and Firewall Management Center event generation.

For each managed device that you want to support Threat Intelligence Director, perform the steps below to configure the associated access control policy.

Elements that are configured to use Threat Intelligence Director after data has been published will automatically receive all currently-published observables.

Procedure

Step 1 Verify that the **Enable Threat Intelligence Director** check box is checked in **General Settings** of the access control policy. To navigate to **General Settings**, choose **Policies** > **Access Control heading** > **Access Control**, and click **Edit** > **More** > **Advanced Settings**. This option is enabled by default.

For more information, see Access Control Policy Advanced Settings, on page 1756.

Step 2 Add rules that allow (rather than trust) connections to the access control policy if they are not already present. Threat Intelligence Director requires that the access control policy specify at least one rule.

Because Threat Intelligence Director depends on inspection, ensure that you allow traffic, rather than trust it, because the purpose of trusting traffic is to bypass inspection. For more information, see Creating a Basic Access Control Policy, on page 1748.

Step 3 If you choose **Intrusion Prevention** as the default action for the access control policy and you want to decrypt traffic for TID detection, associate an SSL policy with the access control policy; see Associating Other Policies with Access Control, on page 1762.

Step 4 If you want `SHA-256` observables to generate observations and Secure Firewall Management Center events:

a) Create a file policy containing one or more **Malware Cloud Lookup** or **Block Malware** file rules.

For more information, see Configure File Policies, on page 2216.

b) Associate this file policy with one or more rules in the access control policy.

Step 5 If you want `IPv4`, `IPv6`, `URL`, or `Domain Name` observations to generate connection and security intelligence events, enable connection and security intelligence logging in the access control policy:

a) In access control rules where you invoked a file policy, enable **Log at End of Connection** and **File Events: Log Files**, if not already enabled.

For more information, see *Logging Connections with Access Control Rules* in the Cisco Secure Firewall Management Center Administration Guide.

b) Verify that default logging (**DNS Policy**, **Networks**, and **URLs**) is enabled in your Security Intelligence settings.

For more information, see *Logging Connections with Security Intelligence* in the Cisco Secure Firewall Management Center Administration Guide.

Step 6 Deploy configuration changes; see Deploy Configuration Changes, on page 259.

What to do next

Complete remaining items in How To Set Up Threat Intelligence Director, on page 2740

Options for Ingesting Data Sources

Choose a configuration option based on the data type and delivery mechanism you want to use.

For more information about these data types, see Source Requirements, on page 2739.

Table 146: Options for Ingesting Data Sources

Data Type	Ingestion Options
STIX	• Ingest STIX feeds from a TAXII server: See Fetch TAXII Feeds to Use as Sources, on page 2743 • Download STIX data from a URL: See Fetch Sources from a URL, on page 2744 • Upload a STIX file: See Upload a Local File to Use as a Source, on page 2745

Data Type	Ingestion Options
Flat file	• Download data from a URL: See Fetch Sources from a URL, on page 2744 • Upload a flat file: See Upload a Local File to Use as a Source, on page 2745

Fetch TAXII Feeds to Use as Sources

If you encounter an issue during TID configuration or operation, see Troubleshoot Threat Intelligence Director, on page 2776

Procedure

Step 1 Make sure your source meets the requirements in Source Requirements, on page 2739

Step 2 Choose **Integration** > **Intelligence** > **Sources**.

Step 3 Click **Add** (+).

Step 4 Choose `TAXII` as the **Delivery** method for the source.

Step 5 Enter information.

- If the host server requires an encrypted connection, configure the **SSL Settings** as described in Configure TLS/SSL Settings for a Threat Intelligence Director Source, on page 2747.
- You cannot change the **Action** selection for TAXII sources.

  `Block` is not an **Action** option for TAXII sources, as STIX data can contain complex indicators, which the system cannot block. Devices (elements) store and take action based on single observables; they cannot take action based on multiple observables.

  However, after ingestion, you can block individual observables and simple indicators obtained from the source. For more information, see Edit Threat Intelligence Director Actions at the Source, Indicator, or Observable Level, on page 2772.
- It may take some time for the list of feeds to load.
- The **Update Every** interval specifies the frequency that Threat Intelligence Director retrieves updates from the TAXII source.

  Set an update frequency that makes sense for how often the data source is updated. For example, if the source is updated 3 times per day, set your update interval to 1440/3 or 480 minutes to regularly capture the latest data.
- After the number of days you specify for **TTL**, Threat Intelligence Director deletes:
  - all of the source's indicators that are not included in subsequent source updates.
  - all observables not referenced by a surviving indicator.

Note

If the source contains no updates for the number of days specified for TTL and the source checksum remains unchanged, downloads are treated as feeds with no updates. For observables to receive a new TTL value, the source must contain a few updates.

Step 6 If you want to immediately begin publishing to elements, confirm that the **PUBLISH Slider** ( ) is enabled.

When this option is enabled, the system automatically publishes the initial source data and any subsequent changes.

For details, see Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level, on page 2774.

Step 7 Click **Save**.

What to do next

- TAXII feeds can contain a lot of data, it may take some time for the system to ingest all of the data. To view ingestion status, refresh the Sources page.
- If you see an error for this source, hover over status for details.
- If you are doing initial Threat Intelligence Director configuration, return to How To Set Up Threat Intelligence Director, on page 2740.

Fetch Sources from a URL

Configure a URL source if you want Threat Intelligence Director to fetch files from a host.

If you encounter an issue during TID configuration or operation, see Troubleshoot Threat Intelligence Director, on page 2776

Procedure

Step 1 Make sure your source meets the requirements in Source Requirements, on page 2739

Step 2 Choose **Integration** > **Intelligence** > **Sources**.

Step 3 Click **Add** ( ).

Step 4 Choose `URL` as the **Delivery** method for the source.

Step 5 Complete the form.

- If you are ingesting a flat file, choose a **Type** that describes the data contained within the source.
- If the host server requires an encrypted connection, configure the **SSL Settings** as described in Configure TLS/SSL Settings for a Threat Intelligence Director Source, on page 2747.
- For Name: To simplify sorting and handling of incidents based on Threat Intelligence Director indicators, use a consistent naming scheme across sources. For example, <source>-<type>.

  Including the source name simplifies returning to the source for further information or feedback.

Be sure to enter the name consistently. For example, for a source with IPv4 addresses, you might always use IPV4 (not IPv4 or ipv4 or IP_v4 or IP_V4 or ip-v4 or IP-v4, IP-V4, etc.)

- If you are ingesting a STIX file, `Block` is not an **Action** option, as STIX data can contain complex indicators, which the system cannot block. Devices (elements) store and take action based on single observables; they cannot take action based on multiple observables.

  However, after ingestion, you can block individual observables and simple indicators obtained from the source. For more information, see Edit Threat Intelligence Director Actions at the Source, Indicator, or Observable Level, on page 2772.

- Set an update frequency that makes sense for how often the data source is updated. For example, if the source is updated 3 times per day, set your update interval to 1440/3 or 480 minutes to regularly capture the latest data.

- After the number of days you specify for the **TTL** interval, Threat Intelligence Director deletes:

  - all of the source's indicators that are not included in subsequent source updates.
  - all observables not referenced by a surviving indicator.

Note
If the source contains no updates for the number of days specified for TTL and the source checksum remains unchanged, downloads are treated as feeds with no updates. For observables to receive a new TTL value, the source must contain a few updates.

Step 6 If you want to immediately begin publishing to elements, confirm that the **Publish Slider** ( ) is enabled.

When this option is enabled, the system automatically publishes the initial source data and any subsequent changes.

For details, see Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level, on page 2774.

Step 7 Click **Save**.

What to do next

- To view ingestion status, refresh the Sources page. If you see an error, hover over status for details.
- If you are doing initial Threat Intelligence Director configuration, return to How To Set Up Threat Intelligence Director, on page 2740.

Upload a Local File to Use as a Source

Use this procedure for a one-time manual upload of a local file.

When ingesting a STIX file, Threat Intelligence Director creates a simple or complex indicator from the contents of the STIX file.

When ingesting a flat file, Threat Intelligence Director creates a simple indicator for each observable value in the file.

If you encounter an issue during Threat Intelligence Director configuration or operation, see Troubleshoot Threat Intelligence Director, on page 2776

Procedure

Step 1 Make sure your file meets the requirements in Source Requirements, on page 2739

Step 2 Choose **Integration** > **Intelligence** > **Sources**.

Step 3 Click **Add** (+).

Step 4 Choose `Upload` as the **Delivery** method for the source.

Step 5 Complete the form.

- If you are uploading a flat file, choose a **Type** that describes the data contained within the source.
- For Name: To simplify sorting and handling of incidents based on Threat Intelligence Director indicators, use a consistent naming scheme across sources. For example, <source>-<type>.

  Including the source name simplifies returning to the source for further information or feedback.

  Be sure to enter the name consistently. For example, for a source with IPv4 addresses, you might always use IPV4 (not IPv4 or ipv4 or IP_v4 or IP_V4 or ip-v4 or IP-v4, IP-V4, etc.)
- If you are uploading a STIX file, `Block` is not an **Action** option, because STIX data can contain complex indicators. Devices (elements) store and take action based on single observables; they cannot take action based on multiple observables.

  However, you can block a simple indicator at the indicator or observable level. For more information, see Edit Threat Intelligence Director Actions at the Source, Indicator, or Observable Level, on page 2772.
- After the number of days you specify for the **TTL** interval, Threat Intelligence Director deletes:
  - all of the source's indicators that are not included in a subsequent upload.
  - all observables not referenced by a surviving indicator.

  Note
  If the source contains no updates for the number of days specified for TTL and the source checksum remains unchanged, downloads are treated as feeds with no updates. For observables to receive a new TTL value, the source must contain a few updates.

Step 6 If you want to immediately begin publishing to elements, confirm that the **Publish Slider** (toggle) is enabled.

If you do not publish the source at ingestion, you cannot publish all source indicators at once later; instead, you must publish each observable individually. See Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level, on page 2774.

Step 7 Click **Save**.

What to do next

- To view ingestion status, refresh the Sources page. If you see an error, hover over status for details.
- If you are doing initial Threat Intelligence Director configuration, return to How To Set Up Threat Intelligence Director, on page 2740.

Handling of Duplicate Indicators

If a single indicator is included in multiple sources:

Each instance of the indicator generates an incident, so one encounter with a particular threat may generate multiple incidents.

To avoid future duplicate incidents, pause publishing of all but one of the duplicated indicators. See Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level, on page 2774.

Configure TLS/SSL Settings for a Threat Intelligence Director Source

Configure **SSL Settings** if the host server requires an encrypted connection.

Before you begin

- Begin configuring a `TAXII` or `URL` source, as described in Fetch TAXII Feeds to Use as Sources, on page 2743 or Fetch Sources from a URL, on page 2744.

Procedure

Step 1 In the **Edit Source** dialog box, expand the **SSL Settings** section.

Step 2 If your server certificate is self-signed:

a) Enable **Self-Signed Certificate**.

b) Choose a **SSL Hostname Verification** method.

- `Strict`—Threat Intelligence Director requires the source **URL** to match the hostname provided in the server certificate.

  If the hostname includes a wildcard, TID cannot match more than one subdomain.

- `Browser Compatible`—Threat Intelligence Director requires the source **URL** to match the hostname provided in the server certificate.

  If the hostname includes a wildcard, TID matches all subdomains.

- `Allow All`—Threat Intelligence Director does not require the source **URL** to match the hostname provided in the server certificate.

For example, if `subdomain1.subdomain2.cisco.com` is your source **URL** and `*.cisco.com` is the hostname provided in the server certificate:

- `Strict` hostname verification fails.
- `Browser Compatible` hostname verification succeeds.
- `Allow All` hostname verification ignores the hostname values completely.

c) For **Server Certificate**:

- If you have access to the PEM-encoded self-signed server certificate, open the certificate in a text editor and copy the entire block of text, including the `BEGIN CERTIFICATE` and `END CERTIFICATE` lines. Enter this entire string into the field.

- If you do not have access to the self-signed server certificate, leave the field blank. After you save the source, Threat Intelligence Director retrieves the certificate from the server.

Step 3 If your server requires a user certificate:

a) Enter a **User Certificate**:

Open the PEM-encoded certificate in a text editor and copy the entire block of text, including the `BEGIN CERTIFICATE` and `END CERTIFICATE` lines. Enter this entire string into the field.

b) Enter a **User Private Key**:

Open the private key file in a text editor and copy the entire block of text, including the `BEGIN RSA PRIVATE KEY` and `END RSA PRIVATE KEY` lines. Enter this entire string into the field.

What to do next

- Take note of the certificate's expiration date. You may want to set a calendar reminder to enter a new server certificate after the current certificate expires.
- Continue configuring the source:
  - Fetch TAXII Feeds to Use as Sources, on page 2743
  - Fetch Sources from a URL, on page 2744

User Roles with Threat Intelligence Director Access

You can use Firewall Management Center user accounts to access the Threat Intelligence Director menus and pages:

- Accounts with the **Admin** or **Threat Intelligence Director User** user role.
- Accounts with a custom user role containing the **Intelligence** permission.

In addition, you can use Firewall Management Center user accounts with the **Admin**, **Access Admin**, or **Network Admin** user role to enable or disable Threat Intelligence Director in your access control policies.

For more information about user accounts, see the *Users for the Firewall Management Center* chapter in the Cisco Secure Firewall Management Center Administration Guide.

About Backing Up and Restoring Threat Intelligence Director Data

You can use the Firewall Management Center to back up and restore all of the data needed for Threat Intelligence Director: Element data, security intelligence events, connection events, Threat Intelligence Director configurations, and Threat Intelligence Director data. For more information, see the *Backup/Restore* chapter in the Cisco Secure Firewall Management Center Administration Guide.

Note If you host Threat Intelligence Director on the active Firewall Management Center in a high availability configuration, the system does not synchronize Threat Intelligence Director configurations and Threat Intelligence Director data to the standby Firewall Management Center. We recommend performing regular backups of Threat Intelligence Director data on your active Firewall Management Center so that you can restore the data after failover.

Before you attempt to restore the Threat Intelligence Director data on the active Firewall Management Center, pause synchronization on the active peer. For more information, see *Pausing Communication Between Paired Firepower Management Centers* in the Cisco Secure Firewall Management Center Administration Guide.

Table 147: Threat Intelligence Director-Related Backup and Restore File Contents

Threat Intelligence Director-Related File Contents	Backup Selection	Restore Selection
Element data	**Back Up Configuration**	**Restore Configuration Data**
Secure Firewall Management Center event data	**Back Up Events**	**Restore Event Data**
Threat Intelligence Director configurations and Threat Intelligence Director data	**Back Up Threat Intelligence Director**	**Restore Threat Intelligence Director Data**

Analyze Threat Intelligence Director Incident and Observation Data

To analyze incident and observation data generated by Threat Intelligence Director elements, use the Incidents table and Incident Details page.

Observation and Incident Generation

Threat Intelligence Director generates an incident when the first observable for an indicator is seen in traffic. Simple indicators are fully realized after a single observation. Complex indicators are partially realized until one or more additional observations fulfill their pattern. Complex indicators need not necessarily be fulfilled during a single transaction; each observable can be fulfilled separately over time, by different transactions.

Note When evaluating an indicator's pattern, Threat Intelligence Director ignores unsupported and invalid objects and observables on the Do Not Block list.

After an incident is fully realized, subsequent observations trigger new incidents.

Figure 452: Example: Indicator Patterns

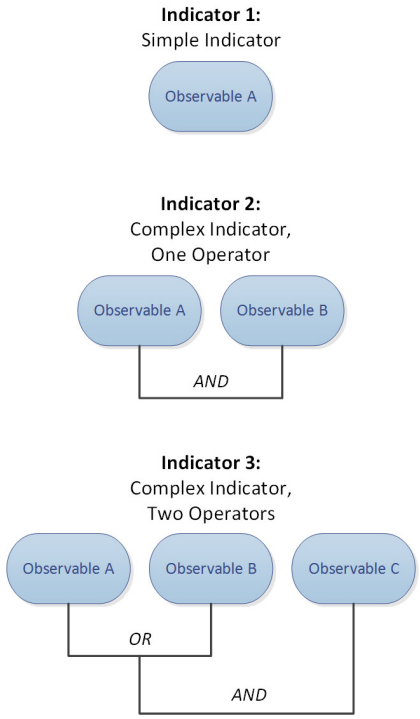

If Threat Intelligence Director ingested the observables from the example above and the observables were seen in order, incident generation would proceed as follows:

1. When the system identifies Observable A in traffic, Threat Intelligence Director:
   - Generates a fully-realized incident for Indicator 1.
   - Generates partially-realized incidents for Indicator 2 and Indicator 3.
2. When the system identifies Observable B in traffic, Threat Intelligence Director:
   - Updates the incident to fully-realized for Indicator 2, since the pattern was fulfilled.
   - Updates the incident to partially-realized for Indicator 3.
3. When the system identifies Observable C in traffic, Threat Intelligence Director:
   - Updates the incident to fully-realized for Indicator 3, since the pattern was fulfilled.
4. When the system identifies Observable A for a second time, Threat Intelligence Director:
   - Generates a new fully-realized incident for Indicator 1.
   - Generates new partially-realized incidents for Indicator 2 and Indicator 3.

If a particular indicator exists in multiple sources, you may see duplicate incidents. For more information, see Troubleshoot Threat Intelligence Director, on page 2776.

Note that incidents are generated only by actual traffic. If there is an observable for URL B, and a user visits URL A which displays a link to URL B, no incident occurs unless the user clicks the URL B link.

View and Manage Incidents

The Incidents page displays summary information for up to 1.1 million of the most recent Threat Intelligence Director incidents; see Incident Summary Information, on page 2752.

Before you begin

- Configure the feature as described in How To Set Up Threat Intelligence Director, on page 2740.
- Understand observation and incident generation, as described in Observation and Incident Generation, on page 2749.

Procedure

Step 1 Choose **Integration** > **Intelligence** > **Incidents**.

Step 2 View your incidents:

- Click **Filter** ( ) to add one or more filters. The default filter is 6 hours. For more information, see Filter Threat Intelligence Director Data in Table Views, on page 2769.
- To view the date and time an incident was last updated by Threat Intelligence Director, hover the cursor over the value in the **Last Updated** column.
- To view more information about the indicator associated with the incident, click the text in the **Indicator Name** column; see View and Manage Indicators, on page 2765.

Step 3 View additional details by clicking a value in the **Incident ID** column.

For an explanation of the details you see, see Incident Details, on page 2752.

- To view indicator details, click an indicator value (for example, an IP address or SHA-256 value) under the **Indicator** heading in the lower section of the window.
- To view observation details, click the arrow to the left of an observation immediately under the **Observations** heading.
- To view this incident on the Security Intelligence Events page, click the **Events** link in the observation details section.

Step 4 (Optional) Enter descriptive information on the incident details page:

Tip: To maximize consistency and usefulness of the options below, plan ahead and document your naming conventions, category choices, and confidence level criteria.

- Enter any value you like in the following fields: **Name**, **Description**, and **Category**.
- Click a rating level for **Confidence**.
- Indicate the status of your investigation into the incident by choosing a value from the drop-down list in the **Status** field.

Incident Summary Information

The Incidents page displays summary information for all Threat Intelligence Director incidents.

Table 148: Incident Summary Information

Field	Description
Last Updated	The number of days since either the system or a user last updated the incident. To view the date and time of the update, hover the cursor over the value in this column.
Incident ID	The unique identifier for the incident. This ID has the following format: `<type>-<date>-<number>` • `<type>`—The type of indicator or observable involved in the incident. For simple indicators, this value indicates the observable type: `IP` (IPv4 or IPv6), `URL` (URL), `DOM` (domain), or `SHA` (SHA-256). For complex indicators, this value is `COM`. • `<date>`—The date (*yyyymmdd*) on which the incident was created. • `<number>`—The daily incident number, that is, a number specifying where the incident occurs in the daily sequence of incidents. Note that this sequence starts at 0. For example, `DOM-20170828-10` is the 11th incident created on that day. Next to the identifier, the system displays an icon that indicates whether the incident is **Partially Realized** or **Fully Realized**. For more information, see Observation and Incident Generation, on page 2749.
Indicator Name	The name of the indicator involved in the incident. To view additional information about the indicator, click the value in this column; see View and Manage Indicators, on page 2765.
Type	The type of indicator involved in the incident. • Indicators that contain a single observable display the data type (`URL`, `SHA-256`, etc.) • Indicators that contain two or more observables display as `Complex`.
Action Taken	The action taken by the system in relation to the incident. For more information, see Incident Details, on page 2752.
Status	The status of your investigation into the incident. For more information, see Incident Details, on page 2752.
Delete ( )	Clicking this icon permanently deletes the incident.

Incident Details

The Incident Details window displays information about a single Threat Intelligence Director incident. This window is divided into two sections:

- Incident Details: Basic Information, on page 2753
- Incident Details: Indicator and Observations, on page 2754

Incident Details: Basic Information

The upper section of the Incident Details window provides the information described below.

Table 149: Basic Incident Information Fields

Field	Description
Partially-Realized *IncidentID* or **Fully-Realized** *IncidentID*	An icon indicating the incident's status (partially-realized or fully-realized), as well as the unique identifier for the incident. **Note** When determining an incident's status, Threat Intelligence Director ignores unsupported and invalid observables and observables on the Do Not Block list.
Opened	The date and time the incident was last updated.
Name	A custom, optional incident name that you enter manually. Tip: If there is information from the source in the Description field (in the bottom part of the window), use information from that field to name the incident.
Description	A custom, optional incident description that you enter manually. Tip: If there is information from the source in the Description field (in the bottom part of the window), use information from that field to describe the incident.
Observations	The number of observations within the incident.
Confidence	An optional rating that you can manually select to indicate the relative importance of the incident.
Action Taken	The action taken by the system: `Monitored`, `Blocked`, or `Partially Blocked`. `Partially Blocked` indicates that the incident contained both `Monitored` and `Blocked` observations. **Note** The **Action Taken** indicates the action taken by the system, not necessarily the action selected in Threat Intelligence Director. For more information, see Threat Intelligence Director-Firewall Management Center Action Prioritization, on page 2757.
Category	A custom, optional tag or keyword that you manually add to the incident.
Status	A value indicating the current stage of your analysis of the incident. All incidents are `New` until you change the **Status** for the first time. This field is optional. Depending on the needs of your organization, consider using the status values as follows: • `New`—The incident requires investigation, but you have not started investigating. • `Open`—You are currently investigating the incident. • `Closed`—You investigated the incident and took action. • `Rejected`—You investigated the incident and determined there was no action to take.
Delete ( )	Clicking this icon permanently deletes this incident.

Incident Details: Indicator and Observations

The lower section of the Incident Details window provides an in-depth view of the indicator and observation information. This information is organized as **Indicator** fields, the indicator pattern, and **Observations** fields.

Indicator Section

When you first view indicator details, this section displays only the indicator name.

Click the indicator name to view the indicator on the Indicators page.

Click the down arrow next to the indicator name to view more indicator details without leaving the incident. Detail fields include:

Table 150: Indicator Fields

Field	Description
Description	The indicator description provided by the source.
Source	The source that contained the indicator. Click this link to access full source details.
Expires	The date and time the incident will expire, based on the source's **TTL** value.
Action	The action associated with the indicator. For more information, see Edit Threat Intelligence Director Actions at the Source, Indicator, or Observable Level, on page 2772.
Publish	The publish setting for the indicator. For more information, see Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level, on page 2774.
Download STIX	If the source type is `STIX`, click this button to download the STIX file.

Indicator Pattern

The indicator pattern is a graphical representation of the observables and operators that comprise the indicator. Operators link the observables within the indicator. AND relationships are indicated with the **AND** operator. OR relationships are indicated with the **OR** operator or by a close grouping of several observables.

If an observable in the pattern has already been seen, the observable box is white. If an observable has not already been seen, the observable box is grey.

In the indicator pattern:

- Click the **Add to Do-Not-Block List** button to add the observable to the Do Not Block list. This icon is present in both white and grey observable boxes. For more information, see About Adding Threat Intelligence Director Observables to the Do Not Block List, on page 2775.
- If you hover the cursor over a white observable box, the system highlights the related observation in the **Observations** section.
- If you click a white observable box, the system highlights the related observation in the **Observations** section, scrolls that observation into view (if multiple observations are present), and expands that observation's detailed display.

- If you hover the cursor over or click a grey observable box in the indicator pattern, there is no change in the **Observations** section. Because the observable is unseen, there are no observation details to display yet.

Observations Section

By default, the **Observations** section displays summary information, which includes:

- The type of observable that triggered the observation (for example, `Domain`)
- The data that comprises the observable
- Whether the observation is the first observation or a subsequent observation (for example, `1st` or `3rd`)

Note If a single observable has been seen three or more times, Threat Intelligence Director displays the first and last observation details. The details for intermediary observations are not available.

- The date and time of the observation
- The action configured for the observable

If you hover the cursor over an observation in the **Observations** section, the system highlights the related observable in the indicator pattern.

If you click an observation in the **Observations** section, the system highlights the related observable(s) in the indicator pattern and scrolls the first related observable into view (if multiple observables are present). Clicking an observation also expands the details of the observation in the **Observations** section.

Observation details include the following fields:

Table 151: Observation Detail Fields

Field	Description
SOURCE	The source IP address and port for the traffic that triggered the observation.
DESTINATION	The destination IP address and port for the traffic that triggered the observation.
ADDITIONAL INFORMATION	DNS and authentication information related to the traffic that triggered the observation.
Events	This clickable link displays if the observation generated connection, security intelligence, file, or malware events. Click the link to view the events in the Secure Firewall Management Center event table; see Cisco Secure Firewall Management Center Administration Guide.

View Events for a Threat Intelligence Director Observation

For more information about the Secure Firewall Management Center events that Threat Intelligence Director observations generate, see Threat Intelligence Director Observations in Secure Firewall Management Center Events, on page 2756.

The system action logged for Threat Intelligence Director-related events can vary, depending on the interaction of Threat Intelligence Director and other Secure Firewall Management Center features. For more information about action prioritization, see Threat Intelligence Director-Firewall Management Center Action Prioritization, on page 2757.

Before you begin

- Configure the feature as described in How To Set Up Threat Intelligence Director, on page 2740.
- Confirm that you enabled event logging required for Threat Intelligence Director in your access control policy, as described in Configure Policies to Support Threat Intelligence Director, on page 2741.

Procedure

Step 1 Choose **Integration** > **Intelligence** > **Incidents**.

Step 2 Click the **Incident ID** value for the incident.

Step 3 Click the observation in the **Indicator** section to display the observation box.

Step 4 Expand the observation box by clicking the arrow in the upper-left corner of the box.

Step 5 Click the **Events** link in the observation information. For more information on the Security Intelligence display, see the *Connection and Security Intelligence Events* chapter in the Cisco Secure Firewall Management Center Administration Guide.

Threat Intelligence Director Observations in Secure Firewall Management Center Events

If you fully configure your access control policy, Threat Intelligence Director observations generate the following Secure Firewall Management Center events:

Table 152: Secure Firewall Management Center Events Generated by Observations

Observation Content	Connection Events Table	Security Intelligence Events Table	File Events Table	Malware Events Table
SHA-256	Yes	No	Yes	Yes, if disposition is Malware or Custom Detection.

Observation Content	Connection Events Table	Security Intelligence Events Table	File Events Table	Malware Events Table
Domain Name, **URL**, or **IPv4/IPv6**	Yes Threat Intelligence Director-related connection events are identified with a Threat Intelligence Director-related **Security Intelligence Category** value.	Yes Threat Intelligence Director-related security intelligence events are identified with a Threat Intelligence Director-related **Security Intelligence Category** value.	No	No

Factors That Affect the Action Taken

Many factors determine when the system takes action and what action the system takes when it detects traffic that matches a Threat Intelligence Director observable.

- Features like Security Intelligence take action before Threat Intelligence Director does. For details, see Threat Intelligence Director-Firewall Management Center Action Prioritization, on page 2757.
- Generally, the action configured for an observable (which may differ from the action configured for its parent indicator or source) is the action that will be taken.
- Because STIX sources can contain complex indicators, the Action setting for the source can be set only to Monitor. However, individual simple indicators or observables contained in a STIX feed or file can be set to Block.
- Action settings for indicators and observables can be inherited or individually configured to override inheritance. See Inheritance in Threat Intelligence Director Configurations, on page 2770 and Edit Threat Intelligence Director Actions at the Source, Indicator, or Observable Level, on page 2772.
- Traffic that might otherwise be actionable might be on a Do Not Block list. For details, see Add Threat Intelligence Director Observables to a Do Not Block List, on page 2776.
- The configured action is taken for both partially- and fully-realized incidents.
- An incident based on a complex indicator can be partially blocked. This can occur if the indicator includes both monitored and blocked observations.
- Pausing publishing affects actions the system takes. See About Pausing Publishing, on page 2772 and Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level, on page 2774.
- Pausing the Threat Intelligence Director feature prevents all actions. After you resume the feature, actionable data may be different from before. For details, see Pause Threat Intelligence Director and Purge Threat Intelligence Director Data from Elements, on page 2773.

Threat Intelligence Director-Firewall Management Center Action Prioritization

If Threat Intelligence Director observable actions conflict with Firewall Management Center policy actions, the system prioritizes actions as follows:

- Security Intelligence Do Not Block
- TID Block
- Security Intelligence Block
- TID Monitor
- Security Intelligence Monitor

Specifically:

Table 153: Threat Intelligence Director URL Observable Action vs. Security Intelligence Action

Setting: Security Intelligence Action	Setting: Threat Intelligence Director Observable Action	Threat Intelligence Director Incidents Field: Action Taken	Security Intelligence Events Fields:		
			Action	Security Intelligence Category	Reason
`Do Not Block`	`Monitor` or `Block`	No TID incident	No Security Intelligence event		
`Block`	`Monitor`	`Blocked`	`Block`	as determined by system analysis; see Security Intelligence Categories, on page 1891	`URL Block`
	`Block`	`Blocked`	`Block`	`TID URL Block`	`URL Block`
`Monitor`	`Monitor`	`Monitored`	Determined by access control rules processed after Security Intelligence and TID.	`TID URL Monitor`	`URL Monitor`
	`Block`	`Blocked`	`Block`	`TID URL Block`	`URL Block`

Table 154: Threat Intelligence Director IPv4/IPv6 Observable Action vs. Security Intelligence Action

Setting: Security Intelligence Action	Setting: Threat Intelligence Director Observable Action	Threat Intelligence Director Incidents Field: Action Taken	Security Intelligence Events Fields:		
			Action	Security Intelligence Category	Reason
`Do Not Block`	`Monitor` or `Block`	No TID incident	No Security Intelligence event		

Setting: Security Intelligence Action	Setting: Threat Intelligence Director Observable Action	Threat Intelligence Director Incidents Field: Action Taken	Security Intelligence Events Fields:		
			Action	Security Intelligence Category	Reason
Block	Monitor	No TID incident	Block	as determined by system analysis; see Security Intelligence Categories, on page 1891	IP Block
	Block	Blocked	Block	TID IPv4 Block TID IPv6 Block	IP Block
Monitor	Monitor	Monitored	Determined by access control rules processed after Security Intelligence and TID.	TID IPv4 Monitor TID IPv6 Monitor	IP Monitor
	Block	Blocked	Block	TID IPv4 Block TID IPv6 Block	IP Block

Table 155: Threat Intelligence Director Domain Name Observable Action vs. DNS Policy Action

Setting: DNS Policy Action	Setting: Threat Intelligence Director Domain Name Observable Action	Threat Intelligence Director Incidents Field: Action Taken	Security Intelligence Events Fields:		
			Action	Security Intelligence Category	Reason
Do Not Block	Monitor or Block	No TID incident	No Security Intelligence event		
Drop, Domain Not Found Sinkhole—Log Sinkhole—Block and Log	Monitor	Blocked	Block	as determined by system analysis; see Security Intelligence Categories, on page 1891	DNS Block
	Block	Blocked	Block	TID Domain Name Block	DNS Block

Setting: DNS Policy Action	Setting: Threat Intelligence Director Domain Name Observable Action	Threat Intelligence Director Incidents Field: Action Taken	Security Intelligence Events Fields:		
			Action	Security Intelligence Category	Reason
Monitor	Monitor	Monitored	Determined by access control rules processed after Security Intelligence and TID.	TID Domain Name Monitor	DNS Monitor
	Block	Blocked	Block	TID Domain Name Block	DNS Block

Table 156: TID SHA-256 Observable Action vs. Malware Cloud Lookup File Policy

File Disposition	Threat Intelligence Director SHA-256 Observable Action	Action Taken in Threat Intelligence Director Incidents	Action in File Events	Action in Malware Events
Clean	Monitor or Block	Monitored	Malware Cloud Lookup	n/a
Malware	Monitor or Block	Monitored	Malware Cloud Lookup	n/a
Custom	Monitor or Block	Monitored	• Malware Cloud Lookup, if SHA-256 is not in a custom detection list. • Custom Detection, if SHA-256 is in a custom detection list.	• Malware Cloud Lookup, if SHA-256 is not in a custom detection list. • Custom Detection, if SHA-256 is in a custom detection list.
Unknown	Monitor or Block	Monitored	Malware Cloud Lookup	n/a

Note Threat Intelligence Director matching occurs before the system sends a file for dynamic analysis.

Table 157: TID SHA-256 Observable Action vs. Block Malware File Policy

File Disposition	Threat Intelligence Director SHA-256 Observable Action	Action Taken in Threat Intelligence Director Incidents	Action in File Events	Action in Malware Events
Clean or Unknown	Monitor	Monitored	Malware Cloud Lookup	n/a
	Block	Blocked	• TID Block, if SHA-256 is not in a custom detection list. Modified file disposition is Custom. • Custom Detection Block, if SHA-256 is in a custom detection list.	TID Block Modified file disposition is Custom.
Malware or Custom	Monitor	Blocked	Block Malware	Block Malware
	Block	Blocked	• TID Block, if SHA-256 is not in a custom detection list. Modified file disposition is Custom. • Custom Detection Block, if SHA-256 is in a custom detection list.	TID Block Modified file disposition is Custom.

View and Change Threat Intelligence Director Configurations

Use the following information to review and fine-tune your configuration as needed.

View Threat Intelligence Director Status of Elements (Managed Devices)

All devices that are registered to the Firewall Management Center as managed devices appear automatically on the Elements page. All properly-configured elements (as specified in Configure Policies to Support Threat

Intelligence Director, on page 2741) will receive all currently-published observables, including those ingested before the element was added.

Procedure

Step 1 Choose **Integration** > **Intelligence** > **Elements**.

Step 2 To see whether the element is connected and Threat Intelligence Director is enabled, hover over the icon beside the element name.

Note
After deploying, it may take up to 5 minutes for information on this page to update, including the applied access control policy and whether TID is enabled.

View and Manage Sources

The Sources page displays summary information about all configured sources; see Source Summary Information, on page 2763.

Procedure

Step 1 Choose **Integration** > **Intelligence** > **Sources**.

Step 2 View your sources:

- To filter the sources displayed on the page, click **Filter** ( ). For more information, see Filter Threat Intelligence Director Data in Table Views, on page 2769.
- To view detailed ingestion status, hover the cursor over the text in the **Status** column. For more information, see Source Status Details, on page 2764.

Step 3 Manage your sources:

- To edit the **Action** setting, see Edit Threat Intelligence Director Actions at the Source, Indicator, or Observable Level, on page 2772. If an action is fixed, it is the only supported action for the source **Type**.
- To edit the **Publish** setting, click **Slider** ( ). For more information, see Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level, on page 2774.
- To pause or resume Threat Intelligence Director updating the source, click **Pause Updates** or **Resume Updates**. If you pause updates, updating is paused but existing indicators and observables remain in TID.

- To delete the source, click **Delete** ( ). Delete is greyed out if the source is still processing. Deleting a source deletes all indicators associated with that source. Associated observables may also be deleted; they are retained if they are associated with indicators remaining in the system.

Source Summary Information

The Sources page displays summary information for all configured sources. The table below provides brief descriptions of the fields in the summary display. For detailed information on these fields, see descriptions in the relevant configuration topic for the source: See Options for Ingesting Data Sources, on page 2742.

Table 158: Sources Summary Information

Field	Description
Name	The source name.
Type	The data format of the source (`STIX` or `Flat File`).
Delivery	The method Threat Intelligence Director uses to retrieve the source.
Action	The action (`Block` or `Monitor`) that the system is configured to perform on traffic matching the data contained within this source. For more information about Threat Intelligence Director actions, including availability, inheritance, and overriding inheritance, see Factors That Affect the Action Taken, on page 2757.
Publish	`On` or `Off` toggle specifying whether Threat Intelligence Director publishes data from the source to registered elements (managed devices configured to support Threat Intelligence Director). Indicators can inherit **Publish** settings from a parent source, and observables can inherit **Publish** settings from a parent indicator. For more information, see Inheritance in Threat Intelligence Director Configurations, on page 2770.
Last Updated	The date and time Threat Intelligence Director last updated the source.

Field	Description
Status	The current status of the source: • **New**—The source is newly created. • **Scheduled**—The initial download or subsequent update is scheduled, but not yet in progress. • **Downloading**—Threat Intelligence Director is performing the initial download or update refresh. • **Parsing** or **Processing**—Threat Intelligence Director is ingesting the source. • **Completed**—Threat Intelligence Director finished ingesting the source. • **Completed with Errors**—Threat Intelligence Director finished ingesting the source, but some observables are unsupported or invalid. • **Error**—Threat Intelligence Director experienced a problem. If the source is a TAXII or URL source with an **Update Frequency** specified, and updates are not paused, Threat Intelligence Director retries on the next scheduled update. Refresh the page to update the status.
Edit ( )	Clicking this icon allows you to edit settings for the source.
Delete ( )	Clicking this icon permanently deletes the source.

Source Status Details

When you hover over a source's **Status** value in the Sources summary page, Threat Intelligence Director provides the additional details described below.

Data	Description
Status Message	Briefly describes the current status of the source.
Last Updated	Specifies the date and time Threat Intelligence Director last updated the source.
Next Update	For TAXII and URL sources, this value specifies when Threat Intelligence Director will update the source next.

Data	Description
Indicators	Specifies indicator counts: • **Consumed**—The number of indicators Threat Intelligence Director processed during the most recent source update. This number represents all indicators contained in the update, regardless of whether they were ingested or discarded. • **Discarded**—The number of malformed indicators that the system did not add to Threat Intelligence Director during the most recent update. **Note** For `TAXII` sources, Threat Intelligence Director provides separate **Last Update** and **Total** indicator counts, because TAXII updates add incremental data, rather than replacing existing data. For indicators from other source types, Threat Intelligence Director provides only the **Last Update** count, because updates from those sources replace the existing data set entirely. If all of an indicator's observables are **Invalid**, Threat Intelligence Director discards the indicator.
Observables	Specifies observable counts: • **Consumed**—The number of observables Threat Intelligence Director processed during the most recent source update. This number represents all observables contained in the update, regardless of whether they were ingested or discarded. • **Unsupported**—The number of unsupported observables that the system did not add to Threat Intelligence Director during the most recent update. For more information about supported observable types, see information about content types in Source Requirements, on page 2739. • **Invalid**—The number of invalid observables that the system did not add to Threat Intelligence Director during the most recent update. An observable is invalid if it is improperly constructed. For example, `10.10.10.10.123` is not a valid IPv4 address. **Note** For `TAXII` sources, Threat Intelligence Director provides separate **Last Update** and **Total** observable counts, because TAXII updates add incremental data, rather than replacing existing data. For observables from other source types, Threat Intelligence Director provides only the **Last Update** count, because updates from those sources replace the existing data set entirely.

View and Manage Indicators

Indicators are generated automatically from ingested sources. For more information about information on this page, see Indicator Summary Information, on page 2766.

Procedure

Step 1 Choose **Integration** > **Intelligence** > **Sources**.

Step 2 Click **Indicators**.

Step 3 View your current indicators:

- To filter the indicators displayed on the page, click **Filter** (). For more information, see Filter Threat Intelligence Director Data in Table Views, on page 2769.
- To view additional details about an indicator (including associated observables), click the indicator name. For more information, see Indicator Details, on page 2767.
- In the **Incidents** column, click the number to view information about incidents associated with an indicator, or hover the cursor over Incidents to view whether the incidents are fully- or partially-realized.
- To determine whether Threat Intelligence Director finished ingesting an indicator from the source, view the **Status** column.

Step 4 Manage your current indicators:

- To edit the **Action**, see Edit Threat Intelligence Director Actions at the Source, Indicator, or Observable Level, on page 2772. If an action is fixed, it is the only supported action for the source **Type**.
- To edit the **Publish** setting, see Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level, on page 2774.
- To add one or more of an indicator's observables to the Do Not Block list, click the indicator name to access the Indicator Details page. For more information, see About Adding Threat Intelligence Director Observables to the Do Not Block List, on page 2775.

Indicator Summary Information

The Indicators page displays summary information for all indicators associated with configured sources.

Table 159: Indicators Summary Information

Field	Description
Type	• Indicators that have a single observable list the data type of that observable (URL, SHA-256, etc.) • Indicators that have two or more observables are listed as `Complex`. Hover over the type to see the specific observable.
Name	The indicator name.
Source	The source that contained the indicator (the parent source).
Incidents	Information about any incidents associated with the indicator: • an icon specifying whether the incident is **Partially** or **Fully** realized • the number of incidents associated with the indicator

Field	Description
Action	The action associated with the indicator. For more information, see Edit Threat Intelligence Director Actions at the Source, Indicator, or Observable Level, on page 2772. Indicators can inherit **Action** settings from a parent source, and observables can inherit **Action** settings from a parent indicator. For more information, see Inheritance in Threat Intelligence Director Configurations, on page 2770.
Publish	The publish setting for the indicator. For more information, see Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level, on page 2774. Indicators can inherit **Publish** settings from a parent source, and observables can inherit **Publish** settings from a parent indicator. For more information, see Inheritance in Threat Intelligence Director Configurations, on page 2770.
Last Updated	The date and time Threat Intelligence Director last updated the indicator.
Status	The current status of the indicator: • **Pending**—Threat Intelligence Director is ingesting the indicator's observables. • **Completed**—Threat Intelligence Director successfully ingested all of the indicator's observables. • **Completed With Errors**—Threat Intelligence Director finished ingesting the indicator, but some observables are unsupported or invalid.

Indicator Details

The Indicator Details page displays indicator and observable data for an incident.

Table 160: Indicator Details Information

Field	Description
Name	The indicator name.
Description	The indicator description provided by the source.
Source	The source that contained the indicator.
Expires	The date and time the indicator will expire, based on the source's **TTL** value.
Action	The action associated with the indicator. For more information, see Edit Threat Intelligence Director Actions at the Source, Indicator, or Observable Level, on page 2772. Indicators can inherit the **Action** setting from a parent source, and observables can inherit the **Action** setting from a parent indicator. For more information, see Inheritance in Threat Intelligence Director Configurations, on page 2770.

Field	Description
Publish	The publish setting for the indicator. For more information, see Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level, on page 2774. Indicators can inherit the **Publish** setting from a parent source, and observables can inherit the **Publish** setting from a parent indicator. For more information, see Inheritance in Threat Intelligence Director Configurations, on page 2770.
Indicator Pattern	The observables and operators that form the indicator's pattern. Operators link the observables within the indicator. AND relationships are indicated with the **AND** operator. OR relationships are indicated with the **OR** operator or by a close grouping of several observables. Optionally, click the **Add to Do-Not-Block List** button to add an observable to the Do Not Block list. For more information, see About Adding Threat Intelligence Director Observables to the Do Not Block List, on page 2775.

View and Manage Observables

The Observables page displays all successfully ingested observables; see Observable Summary Information, on page 2769.

Before you begin

- Configure one or more sources as described in Fetch TAXII Feeds to Use as Sources, on page 2743, Fetch Sources from a URL, on page 2744, or Upload a Local File to Use as a Source, on page 2745.

Procedure

Step 1 Choose **Integration** > **Intelligence** > **Sources**.

Step 2 Click **Observables**.

Step 3 View your current observables:

- To filter the observables displayed on the page, click **Filter** (). For more information, see Filter Threat Intelligence Director Data in Table Views, on page 2769.
- If the information in the **Value** column is cut off, hover over the value.
- To view indicators that contain the observable, click the number in the **Indicators** column. The Incidents page opens with the observable value as the filter. For more information, see View and Manage Indicators, on page 2765.

Step 4 Manage your current observables:

- To edit the **Action**, see Edit Threat Intelligence Director Actions at the Source, Indicator, or Observable Level, on page 2772.
- To edit an observable's **Publish** setting, see Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level, on page 2774.

- To change an observable's expiration date, modify the **TTL** for the parent source. For more information, see View and Manage Sources, on page 2762.
- To add an observable to the Do Not Block list, click the **Add to Do-Not-Block List** button. For more information, see About Adding Threat Intelligence Director Observables to the Do Not Block List, on page 2775.

Observable Summary Information

The Observables page displays summary information for all ingested observables.

Table 161: Observables Summary Information

Field	Description
Type	The type of observable data: `SHA-256`, `Domain`, `URL`, `IPv4`, or `IPv6`.
Value	The data that comprises the observable.
Indicators	The number of parent indicators containing the observable.
Action	The action configured for the observable. For more information, see Edit Threat Intelligence Director Actions at the Source, Indicator, or Observable Level, on page 2772. Indicators can inherit **Action** settings from a parent source, and observables can inherit **Action** settings from a parent indicator. For more information, see Inheritance in Threat Intelligence Director Configurations, on page 2770.
Publish	The publish setting for the observable; see Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level, on page 2774. Indicators can inherit **Publish** settings from a parent source, and observables can inherit **Publish** settings from a parent indicator. For more information, see Inheritance in Threat Intelligence Director Configurations, on page 2770.
Updated At	The date and time Threat Intelligence Director last updated the observable.
Expires	The date that the observable will be automatically purged from Threat Intelligence Director based on **TTL** for the parent indicator.
Add to Do-Not-Block List button	Clicking this button adds the observable to the Do Not Block list; see About Adding Threat Intelligence Director Observables to the Do Not Block List, on page 2775.

Filter Threat Intelligence Director Data in Table Views

Procedure

Step 1 Choose one of the following Threat Intelligence Director table views:

- **Integration** > **Intelligence** > **Incidents**
- **Integration** > **Intelligence** > **Sources**
- **Integration** > **Intelligence** > **Sources**, and click **Indicators**
- **Integration** > **Intelligence** > **Sources**, and click **Observables**

Step 2 Click **Filter** ( ) and choose a filter attribute.

Step 3 Choose or enter a value for that filter attribute.

Filters are case-sensitive.

Step 4 (Optional) To filter by multiple attributes, click **Filter** ( ) and repeat Step 2 and Step 3.

Step 5 To cancel the changes you have made since you last applied the filter, click **Cancel**.

Step 6 Click **Apply** to refresh the table with the filter applied.

Step 7 To remove a filter attribute individually, click **Remove** ( ) next to the filter attribute and click **Apply** to refresh the table.

Inheritance in Threat Intelligence Director Configurations

When Threat Intelligence Director ingests intelligence data from a source, it creates indicators and observables as child objects of that source. On creation, these child objects inherit **Action** and **Publish** settings from the parent configuration.

An indicator inherits these settings from the parent source. An indicator can only have one parent source.

An observable inherits these settings from the parent indicator(s). An observable can have multiple parent indicators.

For more information, see:

- Inheritance of TID Settings from Multiple Parents, on page 2770
- About Overriding Inherited TID Settings, on page 2771

Inheritance of TID Settings from Multiple Parents

If an observable has multiple parent indicators, the system compares the inherited settings from all the parents and assigns the most secure option to the observable. Thus:

- **Action**: `Block` is more secure than `Monitor`
- **Publish**: `On` is more secure than `Off`

For example, SourceA might contribute IndicatorA and related ObservableA:

Setting	SourceA	IndicatorA	ObservableA
Action	`Block`	`Block`	`Block`
Publish	`Off`	`Off`	`Off`

If SourceB later contributes IndicatorB, which also includes ObservableA, the system modifies ObservableA as follows:

Setting	SourceB	IndicatorB	ObservableA
Action	`Monitor`	`Monitor`	`Block` (inherited from IndicatorA)
Publish	`On`	`On`	`On` (inherited from IndicatorB)

In this example, ObservableA has two parents: one parent for its **Action** setting and one parent for its **Publish** setting. If you manually edit the settings for the observable and then revert the settings, the system sets the **Action** setting to the IndicatorA value and the **Publish** setting to the IndicatorB value.

About Overriding Inherited TID Settings

To override an inherited setting, change the setting at the child level; see Edit Threat Intelligence Director Actions at the Source, Indicator, or Observable Level, on page 2772 and Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level, on page 2774. After you override an inherited setting, the child object retains that setting despite changes to the parent object(s).

For example, you might start with the following original settings, with no overrides set:

Setting	SourceA	IndicatorA	ObservableA1	ObservableA2
Publish	`Off`	`Off`	`Off`	`Off`

If you override the setting for IndicatorA, the settings would be the following:

Setting	SourceA	IndicatorA	ObservableA1	ObservableA2
Publish	`Off`	`On`	`On`	`On`

In this case, any changes to the **Publish** setting for SourceA no longer cascade automatically to IndicatorA. However, inheritance from IndicatorA to ObservableA1 and ObservableA2 continues, because the observable settings are not currently set to override values.

If you later override the setting for ObservableA1:

Setting	SourceA	IndicatorA	ObservableA1	ObservableA2
Publish	`Off`	`On`	`Off`	`On`

Any changes to the **Publish** setting for IndicatorA no longer cascade automatically to ObservableA1. However, those changes continue to cascade to ObservableA2, because it is not set to an override value.

At the observable level, you can revert from an override setting to the inherited setting, and the system resumes cascading setting changes automatically from the parent indicator to that observable.

Edit Threat Intelligence Director Actions at the Source, Indicator, or Observable Level

Note:

- Editing the action for a parent sets the action for all children. If you edit the action at the source level, you set the action for all its indicators. If you edit the action at the indicator level, you set the action for all of its observables.
- Editing the action for a child interrupts inheritance. If you edit the action at the indicator level, and subsequently edit it at the source level, the indicator's action is retained until you edit the action for the individual indicator. If you edit the action at the observable level, and subsequently edit it at the indicator level, the observable's action is retained until you edit the action for the individual observable. At the observable level, you can revert automatically to the parent indicator's action. For more information about inheritance, see Inheritance in Threat Intelligence Director Configurations, on page 2770.

You may also want to review other Factors That Affect the Action Taken, on page 2757.

Procedure

Step 1 Choose any of the following:

- **Integration** > **Intelligence** > **Sources**

  Note
  Threat Intelligence Director does not support blocking TAXII sources at the source level. If the TAXII source contains a simple indicator, you can block at the indicator or observable level.

- **Integration** > **Intelligence** > **Sources**, and click **Indicators**

  Note
  Threat Intelligence Director does not support blocking complex indicators. Instead, block individual observables within the complex indicator.

- **Integration** > **Intelligence** > **Sources**, and click **Observables**

Step 2 Use the **Action** dropdown to choose **Monitor** ( ) or **Block** ( ).

Step 3 (Observables only) If you want to resume inheriting the action setting from the parent indicator, click **Revert** next to the **Action** setting for the observable.

About Pausing Publishing

- If you pause publishing at the feature level, the system purges all Threat Intelligence Director observables stored on your elements. This means that Threat Intelligence Director cannot detect, monitor or block threats. Other security features on your system are not affected.
- If you pause publishing at the source, indicator, or observable level, the system removes the paused Threat Intelligence Director observables from your elements, preventing them from matching traffic.

- Pausing publication for a parent pauses all children. If you pause publishing at the source level, you pause publishing for all its indicators. If you pause publishing at the indicator level, you pause publishing for all of its observables.
- Pausing publication for a child interrupts inheritance. If you pause publishing at the indicator level, and subsequently publish at the source level, publishing for the indicator remains paused until you change the individual setting for the indicator. If you pause publishing at the observable level, and subsequently publish at the indicator level, publishing for the observable remains paused until you change the individual setting for the observable. At the observable level, you can revert automatically to the parent indicator's publishing status. For more information about inheritance, see Inheritance in Threat Intelligence Director Configurations, on page 2770.
- Publishing for Uploaded sources can only be paused at the indicator level.
- For a comparison of pausing publishing for an observable vs adding the observable to the Do Not Block list, see About Adding Threat Intelligence Director Observables to the Do Not Block List, on page 2775.
- If you have specified a publish/pause setting for an individual observable or indicator, source updates do not change that setting if the update contains the same observable or indicator.
- Publishing can be disabled on the object management pages. See Modify the Observable Publication Frequency, on page 2775.
- The option on the Sources page to pause updates is not related to publishing data to elements; it applies to updating sources on the Firewall Management Center from feeds.

Pause Threat Intelligence Director and Purge Threat Intelligence Director Data from Elements

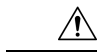

Caution This setting pauses publishing to all elements, purges all Threat Intelligence Director observables stored on your elements, and stops inspecting traffic using the Threat Intelligence Director feature.

To disable observables at a more granular level, see Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level, on page 2774.

Data on the management center (existing incidents and configured sources, indicators, and observables, and ingestion of sources) is not affected by this setting.

Procedure

Step 1 Choose **Integration** > **Intelligence** > **Settings**.

Step 2 Click **Pause**.

What to do next

When you are ready to resume synchronizing Threat Intelligence Director data on your elements and generating observations, manually **Resume** publishing from this page. Existing observables on the management center are published to all elements.

Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level

If publishing is enabled at the Source level, the system automatically publishes the initial source data and any subsequent changes including:

- changes from periodic source refreshes
- changes resulting from system action (for example, **TTL** expiration)
- any user-initiated changes (for example, a change in the **Action** setting for an indicator or observable)

Note To purge all Threat Intelligence Director observables at once from your devices (elements), see Pause Threat Intelligence Director and Purge Threat Intelligence Director Data from Elements, on page 2773.

Before you begin

Before pausing publishing, understand the ramifications described in About Pausing Publishing, on page 2772.

Procedure

Step 1 Choose any of the following:

- **Integration** > **Intelligence** > **Sources**
- **Integration** > **Intelligence** > **Sources**, and click **Indicators**
- **Integration** > **Intelligence** > **Sources**, and click **Observables**

Step 2 Locate the **Publish Slider** ( ) and use it to toggle publishing to elements.

Step 3 (Observables only) If you want to resume inheriting the publication setting from the parent indicator, click **Revert** next to the **Publish** setting for the observable.

What to do next

- Wait at least 10 minutes for elements to receive changes. Changes involving large sources will take longer.
- (Optional) Change the publication frequency for TID data at the observable level; see Modify the Observable Publication Frequency, on page 2775.

Modify the Observable Publication Frequency

By default, the system publishes observables to TID elements every 5 minutes. Use this procedure to set this interval to a different value.

Before you begin

- Enable publication of TID data at the observable level; see Pause or Publish Threat Intelligence Director Data at the Source, Indicator, or Observable Level, on page 2774.

Procedure

Step 1 Choose **Objects** > **Object Management** > **Security Intelligence** > **Network Lists and Feeds**.

Step 2 Click **Edit** (✎) next to the **Cisco-TID-Feed**.

Step 3 Choose a value from the **Update Frequency** drop-down list:

- Choose **Disable** to stop publication of observable data to elements.
- Choose any other value to set the interval for observable publication.

Step 4 Click **Save**.

About Adding Threat Intelligence Director Observables to the Do Not Block List

If you want to exempt an observable in a simple indicator from the specified **Action** (let the traffic pass without monitoring or blocking), you can add the observable to a Do Not Block list.

In a complex indicator, Threat Intelligence Director ignores observables on the Do Not Block list when evaluating traffic, but other observables in that indicator are still evaluated. For example, if an indicator includes Observable 1 and Observable 2 linked by the AND operator, and you add Observable 1 to a Do Not Block list, Threat Intelligence Director generates a fully realized incident when Observable 2 is seen.

By comparison, in the same complex indicator, if you disable publishing of Observable 1 instead of adding it to the Do Not Block list, Threat Intelligence Director generates a partially-realized incident when Observable 2 is seen.

Note If you add an observable to the Do Not Block list, this always takes precedence over the **Action** setting, whether the setting in the observable is an inherited or override value.

Source updates do not affect the Do Not Block list setting for individual observables if the update contains the same observable.

Add Threat Intelligence Director Observables to a Do Not Block List

For detailed information about using Do Not Block lists, see About Adding Threat Intelligence Director Observables to the Do Not Block List, on page 2775.

Tip An "Add to Do Not Block List" button ( ) can appear in several places in the web interface. You can add an observable to a Do Not Block list in any of those locations by clicking this button.

Procedure

Step 1 Choose **Integration** > **Intelligence** > **Sources**, and click **Observables**.

Step 2 Navigate to the observable that you want to allow.

Step 3 Click (**Add to Do-Not-Block List**) for that observable.

What to do next

(Optional) If you need to remove an observable from the Do Not Block list, click the button again.

View a STIX Source File

Procedure

Step 1 Choose **Integration** > **Intelligence** > **Sources**, and click **Indicators**.

Step 2 Click the indicator name.

Step 3 Click **Download STIX**.

Step 4 Open the file in a text editor.

Troubleshoot Threat Intelligence Director

The sections below describe possible solutions and mitigations for common Threat Intelligence Director issues.

Fetching or uploading flat file sources generates an error

If the system fails to fetch or upload a flat file source, check that the data in the flat file matches the **Type** column on the **Integration** > **Intelligence** > **Sources** page.

TAXII or URL source update generates an error

If a TAXII or URL source update generates a source status error, check that your Server Certificate is not expired. If the certificate has expired, enter a new **Server Certificate** or delete the existing **Server Certificate** so Threat Intelligence Director can retrieve a new certificate. For more information, see Configure TLS/SSL Settings for a Threat Intelligence Director Source, on page 2747.

"Block" action is not available for an indicator or source, only "Monitor"

You can change the action for individual observables in the indicator or source.

Threat Intelligence Director table views return "No results"

Table views include the **Sources**, **Indicators**, **Observables**, and **Incidents** pages.

If you do not see data in one of the Threat Intelligence Director table views:

- Check your table filter and consider expanding the time window for the **Last Updated** filter attribute; see Filter Threat Intelligence Director Data in Table Views, on page 2769.
- Verify that you correctly configured your sources; see Options for Ingesting Data Sources, on page 2742.
- Verify that you configured your access control policy and related policies to support Threat Intelligence Director; see Configure Policies to Support Threat Intelligence Director, on page 2741. For example, if your SHA-256 observables are not generating observations, verify that your deployed access control policy contains one or more access control rules that invoke a **Malware Cloud Lookup** or **Block Malware** file policy.
- Verify that you deployed the Threat Intelligence Director-supporting access control policy and related policies to your elements; see Deploy Configuration Changes, on page 259.
- Verify that you did not pause Threat Intelligence Director data publication at the feature level; see Pause Threat Intelligence Director and Purge Threat Intelligence Director Data from Elements, on page 2773.

System is experiencing slowness or decreased performance

For more information about performance impact, see Performance Impact of Threat Intelligence Director, on page 2737.

Secure Firewall Management Center table views do not show Threat Intelligence Director data

If you are publishing observables to your elements but no Threat Intelligence Director data appears in the connection, security intelligence, file, or malware events tables, check the access control and file policies deployed to your elements. For more information, see Configure Policies to Support Threat Intelligence Director, on page 2741.

One or more elements are overwhelmed by Threat Intelligence Director data

If Threat Intelligence Director data is overwhelming one or more of your devices, consider pausing Threat Intelligence Director publishing and purging the data stored on your elements. For more information, see Pause Threat Intelligence Director and Purge Threat Intelligence Director Data from Elements, on page 2773.

System is performing a Malware Cloud Lookup instead of a TID block

This is by design. For more information, see Threat Intelligence Director-Firewall Management Center Action Prioritization, on page 2757.

System is performing a Security Intelligence or DNS Policy action instead of a TID action

This is by design. For more information, see Threat Intelligence Director-Firewall Management Center Action Prioritization, on page 2757.

TID is disabled

- Add memory to your appliance. Threat Intelligence Director can only be used on appliances with at least 15GB of memory.
- Enable REST API access for the Secure Firewall Management Center. For more information, see *Enabling REST API Access* in the Cisco Secure Firewall Management Center Administration Guide.

The system does not generate the Threat Intelligence Director incident or take the Threat Intelligence Director action that you expected

- Verify that all of your managed devices are properly enabled and configured for Threat Intelligence Director. See View Threat Intelligence Director Status of Elements (Managed Devices), on page 2761 and Configure Policies to Support Threat Intelligence Director, on page 2741.
- It takes at least 5-10 minutes for changes to be published to elements, and significantly longer if publishing a large data feed.
- Check the action setting for the observable. See View and Manage Observables, on page 2768.
- For a list of the other factors that influence the Threat Intelligence Director action that the system takes, see Factors That Affect the Action Taken, on page 2757.
- Elements (managed devices) may not have the threat data you think they have. See About Pausing Publishing, on page 2772.

One encounter with a particular threat generates multiple incidents

This can occur if a single indicator is included in multiple sources.

For details, see Handling of Duplicate Indicators, on page 2747.

History for Threat Intelligence Director

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Handling of an indicator that is included in multiple STIX feeds	7.1	Any	If STIX feeds contain identical indicators, an indicator is created for each feed, which may lead to multiple incidents being generated for the same indicator. Previously, only the feed that was downloaded last took effect.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details

Change in action prioritization	6.5	Any	These changes apply if more than one Firepower feature could apply to a particular observable.
			TID blocking/monitoring observable actions now have priority over blocking/monitoring by Security Intelligence.
			Important
			The system still effectively handles traffic as before. Traffic that was previously blocked is still blocked, and monitored traffic is still monitored. This simply changes the component reported in the event as responsible for the action. You may also see more TID incidents generated.
			• If you configure the Block TID observable action, even if the traffic also matches a Security Intelligence Block action:
			  • The Security Intelligence category in the connection event is a variant of TID Block.
			  • The system generates a TID incident with an action taken of Blocked.
			• If you configure the Monitor TID observable action, even if the traffic also matches a Security Intelligence Monitor rule:
			  • The Security Intelligence category in the connection event is a variant of TID Monitor
			  • The system generates a TID incident with an action taken of Monitored.
			Previously, in each of these cases, the system reported the category by analysis and did not generate a TID incident.

Secure Firewall Threat Intelligence Director	6.2.2	Any	Feature introduced: Lets you use threat intelligence from external sources to identify and process threats. New screens: A new top-level **Intelligence** menu with multiple tabs. Supported platforms: Secure Firewall Management Center